



XenServer 9

Contents

About XenServer 9	6
What's new	8
Early Access channel updates	12
Normal channel updates	17
Fixed issues	23
Known Issues	23
Deprecation	26
Quick start	28
XenServer® Enterprise Reference Architecture	50
System requirements	72
Configuration limits	77
Hardware drivers	82
Guest operating system support	83
Connectivity requirements	91
Technical overview	97
Technical FAQs	105
Licensing overview	122
Licensing FAQ	129
Install	136
Other installation scenarios	144
Upgrade from XenServer 8.4	157
Update your XenServer® hosts	167
Apply updates by using XenCenter®	171

Apply updates by using the xe CLI	190
XenCenter®	206
Using XenServer® with Citrix® products	207
Upgrade scenarios for XenServer and Citrix Virtual Apps and Desktops™	210
IntelliCache™	225
PVS-Accelerator	231
Resource pools	239
Manage your pools	244
Manage your hosts	256
Certificate verification	268
GFS2 clustered pools	272
Troubleshoot GFS2 clustered pools	280
Manage users	286
Role-based access control	296
RBAC roles and permissions	297
Use RBAC with the CLI	311
Secure Boot for XenServer 9	316
Networking	317
Manage networking	335
Dom0 network interface naming on XenServer 9	357
Troubleshoot networking	359
Storage	363
Create a storage repository	370
Local storage	373

ISO library on NFS	377
ISO library on SMB/CIFS	379
NFS SRs	381
SMB SR	385
Software iSCSI storage	388
Hardware HBA LUNs	392
Thin-provisioned shared GFS2 block storage	397
Manage storage repositories	408
Storage multipathing	417
Storage read caching	424
Graphics overview	427
Prepare host for graphics	431
Create vGPU enabled VMs	437
Memory usage	440
Monitor and manage your deployment	442
Monitor CPU usage	455
NUMA optimization in XenServer 9	459
Alerts	461
Monitor host and dom0 resources with NRPE	472
Monitor host and dom0 resources with SNMP	480
Manage virtual machines	488
Windows VMs	498
XenServer® VM Tools for Windows	510
Linux VMs	532

Migrate VMs	545
Import and export VMs	550
VM placement	567
Delete VMs	572
vApps	574
Manage UEFI Secure Boot certificate remediation	578
VM memory	581
Advanced notes for virtual machines	587
Troubleshoot VM problems	596
High availability	604
Configure high availability	611
Disaster recovery and backup	616
Enable Disaster Recovery	619
vApps	624
Back up and restore hosts and VMs	625
VM snapshots	630
Cope with machine failures	636
Third-party backup solutions	640
Workload Balancing	641
What's new in Workload Balancing	644
Get started with Workload Balancing	645
Workload Balancing basic tasks	659
Configure Workload Balancing behavior	680
Administer Workload Balancing	717

Certificates for Workload Balancing	750
Troubleshoot Workload Balancing	759
Conversion Manager	768
What's new in Conversion Manager	775
Get started with Conversion Manager	775
Troubleshoot Conversion Manager	793
Command-line interface	795
Infrastructure automation	923
Automating XenServer® with Terraform	923
Troubleshooting	927
Advanced troubleshooting	931
Support	936
Third party notices	941
XenServer® Open Source Licensing and Attribution	941
Develop for XenServer®	947
Data Governance	947

About XenServer 9

June 29, 2026

XenServer® is a virtualization platform that allows organizations to create and manage virtualized server infrastructures. It is designed to optimize the delivery of Windows and Linux virtual machines, providing a robust and scalable solution for data center virtualization.

XenServer offers features such as live migration, snapshot and cloning capabilities, and resource pooling, allowing for efficient management of virtualized workloads. It provides a secure and high-performance environment for running applications and services, making it a suitable choice for businesses looking to streamline their server infrastructure.

Cloud Software Group (CSG) emphasizes XenServer's integration with Citrix® products, creating a comprehensive virtualization and application delivery solution. We aim to enhance the flexibility, agility, and cost-effectiveness of IT operations through centralized management and efficient resource utilization. Overall, XenServer is positioned as a reliable virtualization solution for businesses seeking a powerful and versatile platform for their server environments.

[DOWNLOAD XENSERVER](#)

Note:

If you are using a Citrix License Server that hosts license files to license your hosts and pools, you must change to using licenses held within the License Activation Service in Citrix Cloud.

For more information about getting a XenServer license, see the [XenServer website](#). XenServer is now an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more.](#)

Is XenServer 9 for me?

XenServer 9 is the latest release of XenServer. It succeeds XenServer 8.4 and includes many new features. For more information about these features, see [What's New](#).

Why choose XenServer 9?

- You want to try the latest features from XenServer and can regularly update your hosts and pools.

- NUMA Awareness: Enhanced performance through optimizations for Non-Uniform Memory Access (NUMA) server CPU architectures.
- Driver Multi-Versioning: Flexibility to support multiple driver versions without requiring driver disks.
- Host Secure Boot: Strengthened security with UEFI Secure Boot for host systems.

In any of these cases, XenServer 9 is for you.

[Download XenServer 9 here!](#)

XenCenter®

XenServer 9 requires the latest version of XenCenter, which has a version number of the form “XenCenter YYYY.x.x”. Previous versions of XenCenter, such as XenCenter 8.2.x, are not supported with XenServer 9.

- [Download the latest XenCenter](#)
- [See the documentation](#)

XenCenter YYYY.x.x is fully supported with XenServer 9.

Frequent updates across the XenServer 9 lifecycle

With XenServer 9, frequent updates are made available to you in XenCenter, allowing you to benefit from a more efficient release process that delivers new features and bug fixes at a faster cadence than was previously possible. This feature enables you to experience the frequent update model for administering updates to your XenServer pools and hosts.

During its lifecycle, XenServer 9 provides a stream of frequent and easy-to-apply updates, which enable you to consume new features and bug fixes at the earliest possible juncture. You must apply all available updates periodically. As a result, the behavior and feature set in XenServer 9 can change.

Preview features

The frequent updates can include preview features, which give you an early look at capabilities that might be included in the product later. These preview features can be deployed in a production environment, but are not supported for production use. For more information about preview features, see [Support](#).

Get started

Steps to use XenServer 9:

1. Get XenServer 9 from the [XenServer downloads page](#).
2. [Install the latest version of XenCenter](#).
3. [Install or upgrade to XenServer 9](#).
4. [Apply updates by using XenCenter](#).

What's new

June 29, 2026

Our goal is to deliver new features and product updates to XenServer 9 customers as soon as they are ready. New releases provide more value, so there's no reason to delay updates. Through the XenServer 9 release stream, we deliver updates incrementally in waves to help ensure product quality and maximize availability.

Note:

If you are using a Citrix License Server that hosts license files to license your hosts and pools, you must change to using licenses held within the License Activation Service in Citrix Cloud. For more information, see [Licensing](#).

For more information about getting a XenServer license, see the [XenServer website](#). XenServer is now an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more](#).

Licensing changes

File based licensing system EOL notification

License Activation Service (LAS) is the next generation technology for product activations across the suite of Citrix products.

Note:

File based licensing, traditionally used for activating various on-premises components, reached End of Life (EOL) on April 15, 2026.

As of April 15, 2026, LAS is the only way to activate and license Citrix on-premises components. To remain supported, your XenServer deployment must be on a LAS compatible version.

Requirement Should you have questions or concerns, contact Customer Care. Citrix may limit or suspend your Citrix Maintenance for non-compliance with these requirements without liability in addition to any other remedies Citrix may have at law or equity. These requirements don't apply where prohibited by law or regulation.

License Activation Service

Register your Citrix License Server with Citrix Cloud to use the licenses and entitlements that your License Activation Service holds. For more information, see [Licensing](#).

Currently, you can still use the previous licensing mechanism to host license files on your Citrix License Server. However, in future, XenServer can only be licensed through the License Activation Service.

Citrix Virtual Apps and Desktops™ and Citrix DaaS

For more information about getting a XenServer license, see the [XenServer website](#).

XenServer is now an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more](#).

To learn which versions of Citrix Virtual Apps and Desktops (MCS) and Citrix Provisioning (PVS) are supported with XenServer 9, see [Supported Hypervisors for Citrix Virtual Apps and Desktops \(MCS\) and Citrix Provisioning \(PVS\)](#).

Trial Edition

You can try XenServer 9 for free with Trial Edition. The Trial Edition lets you try Premium Edition features for up to 90 days, but in a restricted size pool of up to 3 hosts. After the 90-day trial period expires, you cannot start any VMs on your Trial Edition host or pool. For more information about the different editions of XenServer, see [XenServer editions](#).

Core Enhancements

- NUMA Optimizations
- Host Secure Boot
- Driver Multi-Versioning

- Improved Host Network Name Handling
- Updated Xen Hypervisor (v4.21)
- Dom0 Kernel Upgraded to 6.x
- New base Operating System for the Control Domain
- GFS2 SRs Now Use Corosync3

Note:

GFS2 SRs on XenServer 8.4 must be converted from Corosync2 to Corosync3 before upgrading to XenServer 9. Refer to the Corosync feature documentation.

NUMA optimization XenServer® 9 introduces NUMA optimization to improve VM performance on NUMA hosts. NUMA optimization co-locates VM vCPUs and memory on the same NUMA node to reduce memory-access latency. This feature is enabled by default, persists across VM lifecycle operations and reboots, and adds per-VM RRD metrics for visibility and troubleshooting. Performance gains vary by workload and CPU architecture - for more information, see [NUMA optimization in XenServer 9](#).

Secure Boot for hosts XenServer® 9 adds support for Secure Boot in firmware for host installation and runtime, including boot from SAN. Secure Boot enforces signature verification for each boot component to prevent untrusted code from running and to improve platform integrity. Mixed pools (hosts with Secure Boot enabled and disabled) are supported. You can enable Secure Boot after an upgrade from XenServer 8.4 or on fresh installs. For more information see [Secure Boot for XenServer 9](#).

Predictable dom0 network interface naming XenServer® 9 adopts systemd predictable network interface names in the control domain (dom0) and consolidates interface management in Toolstack networkd. Interfaces are no longer forcibly renamed to the legacy ethN pattern (for example, eth0, eth1). This change improves stability and maintainability, reduces ambiguity across reboots and hardware changes, and aligns naming with industry standards. Update scripts and automation that assume ethN device names to discover interfaces dynamically (for example, by MAC address or by using `xe pif-list`). For more information, see [Dom0 network interface naming on XenServer 9](#).

New base Operating System for the Control Domain The XenServer® 9 control domain (dom0) operating system is now its own XenServer OS distribution. This custom distribution has been tailored for optimum performance and functionality, and has had unnecessary content removed to reduce potential attack surface. For more information, see [Control domain](#).

Changes to guest operating system support

For the full list of supported guest operating systems in XenServer 9, see [Guest operating system support](#).

Removed

Support for the following guests has been removed in XenServer 9:

- SUSE Linux Enterprise Desktop 12 SP4 (64-bit)

Restrict use of port 80

For new installations, XenServer 9 disables HTTP (port 80) on the management interface and requires TLS (HTTPS) for secure connections. Upgrades retain existing port settings.

For more information about how to open port 80, see [Restrict use of port 80](#).

Removed features

The following features are no longer supported in XenServer 9:

- Host BIOS Boot Support
- Legacy NVIDIA GPU Support (Maxwell & Pascal)
- Software FCoE
- Linux Bridge Network Stack

Compatibility notes

XenServer 9 is compatible with the following components:

- The latest version of the XenServer VM Tools for Windows
- The latest version of the XenServer VM Tools for Linux
- The latest version of the Workload Balancing virtual appliance
- The latest version of the Conversion Manager virtual appliance

These components are available on the [XenServer Downloads page](#).

Early Access channel updates

June 19, 2026

The following features, preview features, improvements, and bug fixes are available in the Early Access update channel. Some of the latest listed entries might not be available in the Normal channel yet.

Note:

This article doesn't list all changes in the Early Access channel, just a subset. For the full and up-to-date set of changes available, see the information in the XenCenter® **Updates** view.

File based licensing system EOL notification

License Activation Service (LAS) is the next generation technology for product activations across the suite of Citrix® products.

Note:

File based licensing, traditionally used for activating various on-premises components, reached End of Life (EOL) on April 15, 2026.

As of April 15, 2026, LAS is the only way to activate and license Citrix on-premises components. To remain supported, your XenServer® deployment must be on a LAS compatible version.

Requirement

Should you have questions or concerns, contact Customer Care. Citrix may limit or suspend your Citrix Maintenance for non-compliance with these requirements without liability in addition to any other remedies Citrix may have at law or equity. These requirements don't apply where prohibited by law or regulation.

Jun 19, 2026

Checksum: be94a90c9b2bdd398765fe03a01ee31e71260c61d2441aa702e362ab800ce2eb

These updates contain the following new features:

- Secure Boot for hosts is now fully supported using the standard Microsoft Secure Boot certificates.

These updates contain the following improvements:

- Update kernel to v6.6.138.
- Update intel-i40e driver to v2.28.16.
- Update AMD microcode to the 2026-05-19 drop.
- Update QEMU datapath to use QEMU 10.2.2.

These updates include fixes for the following issues:

- If a license server does not properly respond, an apply-edition operation may block indefinitely.
- Large multi-part tasks in the toolstack may have their expiry incorrectly calculated resulting in premature timeouts.
- Certain update failures may be misreported due to an error parsing the output of a low level tool.
- In rare cases a corrupt or incomplete offline update bundle could be erroneously accepted.
- Specific firmware can place MMCFG regions in an E820 memory hole, resulting in them being inaccessible to the hypervisor, preventing use of SR-IOV devices in VMs.
- Xen may fail to boot on certain systems due to an issue with merged PDX ranges.
- Xen may fail to boot on certain systems due to an issue with non-aligned memory affinity ranges.
- In rare cases on VM shutdown when using a GFS2 or XFS SR a core dump could be generated.

Jun 9, 2026

Checksum: b9b4e2cf71bc9496205e35ecb06ab467df2c345f08684b8ecf2fd4ac8b7a199f

These updates include fixes for the following issues:

- In rare circumstances storage operations on remote hosts can fail and not be recoverable.
- When cleaning up configuration keys on virtual disks following a host crash, some normally short lived keys can get left behind.

Jun 3, 2026

Checksum: d452d686f71c44d580e8eebf38ff428d71e20d05b552a4824abc8639a062173

These updates include fixes for the following issues:

- Some systems with Intel Granite Rapids processors may fail to boot with the IPU 2026.2 microcode.

Jun 2, 2026

Checksum: 37d7122ebbd6c6a496c3dd58c2ac1d867204e38e5065b831b0e6d6df19a3b45a7

These updates contain the following improvements:

- Server side (API/CLI) functionality to enable updating an existing VM's Secure Boot certificates. For more information, see [Manage UEFI Secure Boot Certificate Remediation](#).

These updates include fixes for the following issues:

- After an emergency HA disable, the toolstack on the pool coordinator may not start correctly and require a manual toolstack restart.
- In rare cases where updates require a mandatory host reboot, livepatches were still being applied when not required.
- When running with host Secure Boot enabled, changing a vGPU from full pass-through of the whole GPU to another GPU type can result in a VM failing to start.

These updates contain upstream code changes for the following issues:

- CVE-2026-31431
- CVE-2026-35091
- CVE-2026-35092

May 21, 2026

Checksum: 2efe82f761eb4c025456f998caaca802fd9f6d40b9919356c128a1ca5cc73726

These updates contain the following improvements:

- Update the device model (QEMU) to v10.1.

These updates include fixes for the following issues:

- A change to avoid VM memory overhead double-counting had unexpected side-effects affecting migrations in rare cases.
- Under extremely rare conditions VMs may start without disks attached.

May 12, 2026

Checksum: 7c34a886e1ecf2b3e36d30f088df79515ab9cd9f8ef838169bf9739276a2c822

These updates include security fixes. For more information, see the Security Bulletins [CTX696527](#), [CTX696562](#) and [CTX696581](#).

These updates contain the following improvements:

- Update Xen hypervisor to v4.21
- Update the Intel microcode to the IPU 2026.2 drop.
- Update the intel-ice driver to v2.4.5.
- Add mpt3sas driver.

These updates include fixes for the following issues:

- In some cases AD user login can fail due to a failure to retrieve subject details.
- In some cases attempting to join a freshly installed host to an existing pool may fail with a [POOL_JOINING_SM_FEATURES_INCOMPATIBLE](#) error.

May 6, 2026

Checksum: 9a06e31a2332908fa02d73fd67ba0c65d3bcd3ced1cdfbf21b0f6d4903a2e7e

These updates include fixes for the following issues:

- In rare circumstances following a host crash the crashdump process can ‘hang’ and the machine does not reboot.
- Live patches that were not applicable may be incorrectly reported as applicable on hosts.
- VM memory overhead was double-counted during migration, causing the reported maximum memory to grow.
- The audit log endpoint does not return data if a request is made using the [since](#) parameter.

Apr 29, 2026

Checksum: deb09b5273ee03884c20af32740cd1339becb275f45b03539a69c2faeda5e740

These updates include the following improvements:

- Updates to collected telemetry. For more information, see [Data governance](#).

Apr 28, 2026

Checksum: 418483727f708033ddfd6a6493c47bf9630c185d0c1a9f5c6190ec62a22b0b794

These updates include security fixes. For more information, see the [Security Bulletin](#).

These updates include upstream code changes for the following issues:

- CVE-2026-23557

Apr 16, 2026

Checksum: 2c50b5593077d548d4c7546d8f01083c5b7c42fb1e2b44deca954641ff1ac187

These updates include fixes for the following issues:

- A recent change caused GFS2 and XFS datapaths to use an increased amount of Control Domain memory.
- If a scan on an LVM-based Storage Repository finds a virtual disk in the metadata that is not known to the toolstack it will (re)introduce it. In the event the VDI is not still present on the SR this (re)introduction will encounter an error.
- BIOS boot VMs may fail to network boot if a vlan priority tag is set in the DHCP response.

Apr 6, 2026

Checksum: 25c82a53636a5344189fe29c2fc7c950ecf96cb01bc8cccd4cd3af6605dfedae

These updates include the following improvements:

- Update OpenSSL to v3.5.5
- Improves the ability of a pool using GFS2 SRs to fence specific hosts which are having networking issues, rather than allowing them to negatively affect the whole cluster.
- Updates the Secure Boot Forbidden Signature Database (DBX) provided to new VMs to the latest published version.
- Add broadcom-mpi3mr driver.
- When using a VM with disks on a GFS2 or XFS SR, a serial number is now supplied.

These updates include fixes for the following issues:

- When the License Server is unreachable while a XenServer host starts up, the host may not become enabled.
- Certain records (e.g. PCI data) are left behind in the pool database after ejecting a host from a pool.
- A toolstack restart after an update synchronization may block application of updates until another synchronization is performed.
- When installing a UEFI VM from a local DVD, the VM may hang at the firmware screen before starting installation.
- In rare circumstances, the device model can fail due to a problem managing the Xen map cache.
- After upgrading an HA-enabled GFS2 pool from an older XenServer release, HA may not be able to be re-enabled.
- When using an LIO iSCSI target, multipath failover may be delayed.

Mar 17, 2026

Checksum: 0e3647a4f77c8f52e9df18776988a3df99809565d3231cbdc7b6b68a79710a76

These updates include security fixes. For more information, see the [Security Bulletin](#).

Mar 11, 2026

Checksum: eab08fa137afaa1f005206692ca767087817f9e475cd99aa7836cc07ba1b53c3

These updates include the following improvements:

- Add qllogic-qla2xxx driver.
- Optimizations to VM migration process to reduce migration time for heavily loaded VMs on systems using AMD processors.

These updates include upstream code changes for the following issues:

- CVE-2024-12243
- CVE-2025-32990
- CVE-2025-32988
- CVE-2025-32989
- CVE-2025-6395
- CVE-2025-9820
- CVE-2025-14831
- CVE-2026-1584

Mar 2, 2026

Checksum: fcdc08d17f088b98e6326f8b1173b15803b5135e5759c3f34c3f631157607358

These updates contain foundational changes.

Feb 23, 2026

Checksum: bb6f87ab348e9bb0ce519aac852952f13bf9e59231f67447d41511a77c69dc02

These updates include fixes for the following issues:

- In some situations, a VM export operation can fail to complete.
- A VDI copy operation initiated by a user with the “VM Power Admin” role may incorrectly result in a permissions error.

Normal channel updates

June 22, 2026

Updates progress from the Early Access update channel to Normal on a regular cadence. The following features, preview features, improvements, and bug fixes are available in the Normal update channel.

Note:

This article doesn't list all changes in the Normal channel, just a subset. For the full and up-to-date set of changes available, see the information in the XenCenter® **Updates** view.

File based licensing system EOL notification

License Activation Service (LAS) is the next generation technology for product activations across the suite of Citrix® products.

Note:

File based licensing, traditionally used for activating various on-premises components, reached End of Life (EOL) on April 15, 2026.

As of April 15, 2026, LAS is the only way to activate and license Citrix on-premises components. To remain supported, your XenServer® deployment must be on a LAS compatible version.

Requirement

Should you have questions or concerns, contact Customer Care. Citrix may limit or suspend your Citrix Maintenance for non-compliance with these requirements without liability in addition to any other remedies Citrix may have at law or equity. These requirements don't apply where prohibited by law or regulation.

Jun 22, 2026

Checksum: be94a90c9b2bdd398765fe03a01ee31e71260c61d2441aa702e362ab800ce2eb

These updates contain the following new features:

- Secure Boot for hosts is now fully supported using the standard Microsoft Secure Boot certificates.

These updates contain the following improvements:

- Update kernel to v6.6.138.
- Update intel-i40e driver to v2.28.16.
- Update AMD microcode to the 2026-05-19 drop.

- Update QEMU datapath to use QEMU 10.2.2.

These updates include fixes for the following issues:

- If a license server does not properly respond, an apply-edition operation may block indefinitely.
- Large multi-part tasks in the toolstack may have their expiry incorrectly calculated resulting in premature timeouts.
- Certain update failures may be misreported due to an error parsing the output of a low level tool.
- In rare cases a corrupt or incomplete offline update bundle could be erroneously accepted.
- Specific firmware can place MMCFG regions in an E820 memory hole, resulting in them being inaccessible to the hypervisor, preventing use of SR-IOV devices in VMs.
- Xen may fail to boot on certain systems due to an issue with merged PDX ranges.
- Xen may fail to boot on certain systems due to an issue with non-aligned memory affinity ranges.
- In rare cases on VM shutdown when using a GFS2 or XFS SR a core dump could be generated.

Jun 15, 2026

Checksum: b9b4e2cf71bc9496205e35ecb06ab467df2c345f08684b8ecf2fd4ac8b7a199f

These updates include fixes for the following issues:

- In rare circumstances storage operations on remote hosts can fail and not be recoverable.
- When cleaning up configuration keys on virtual disks following a host crash, some normally short lived keys can get left behind.

Jun 8, 2026

Checksum: d452d686f71c44d580e8eebfe38ff428d71e20d05b552a4824abc8639a062173

These updates contain the following improvements:

- Server side (API/CLI) functionality to enable updating an existing VM's Secure Boot certificates. For more information, see [Manage UEFI Secure Boot Certificate Remediation](#).

These updates include fixes for the following issues:

- After an emergency HA disable, the toolstack on the pool coordinator may not start correctly and require a manual toolstack restart.
- In rare cases where updates require a mandatory host reboot, livepatches were still being applied when not required.
- When running with host Secure Boot enabled, changing a vGPU from full pass-through of the whole GPU to another GPU type can result in a VM failing to start.

These updates contain upstream code changes for the following issues:

- CVE-2026-31431
- CVE-2026-35091
- CVE-2026-35092

Jun 3, 2026

Checksum: 8e04e0445a637b8242540d311fd76b173dbe74cd626bf55676097574197e77a2

These updates include fixes for the following issues:

- Some systems with Intel Granite Rapids processors may fail to boot with the IPU 2026.2 microcode.

May 25, 2026

Checksum: 2efe82f761eb4c025456f998caaca802fd9f6d40b9919356c128a1ca5cc73726

These updates include security fixes. For more information, see the Security Bulletins [CTX696527](#), [CTX696562](#) and [CTX696581](#).

These updates contain the following improvements:

- Update Xen hypervisor to v4.21
- Update the device model (QEMU) to v10.1.
- Update the Intel microcode to the IPU 2026.2 drop.
- Update the intel-ice driver to v2.4.5.
- Add mpt3sas driver.

These updates include fixes for the following issues:

- Under extremely rare conditions VMs may start without disks attached.
- In some cases AD user login can fail due to a failure to retrieve subject details.
- In some cases attempting to join a freshly installed host to an existing pool may fail with a [POOL_JOINING_SM_FEATURES_INCOMPATIBLE](#) error.
- In rare circumstances following a host crash the crashdump process can ‘hang’ and the machine does not reboot.
- Live patches that were not applicable may be incorrectly reported as applicable on hosts.
- The audit log endpoint does not return data if a request is made using the [since](#) parameter.

May 4, 2026

Checksum: deb09b5273ee03884c20af32740cd1339becb275f45b03539a69c2faeda5e740

These updates include the following improvements:

- Updates to collected telemetry. For more information, see [Data governance](#).

Apr 28, 2026

Checksum: f4a08929c601aefabff1522d038496d2b8362fb976a673759f472e9b8ebd3ab0

These updates include security fixes. For more information, see the [Security Bulletin](#).

These updates include upstream code changes for the following issues:

- CVE-2026-23557

Apr 20, 2026

Checksum: 2c50b5593077d548d4c7546d8f01083c5b7c42fb1e2b44deca954641ff1ac187

These updates include fixes for the following issues:

- A recent change caused GFS2 and XFS datapaths to use an increased amount of Control Domain memory.
- If a scan on an LVM-based Storage Repository finds a virtual disk in the metadata that is not known to the toolstack it will (re)introduce it. In the event the VDI is not still present on the SR this (re)introduction will encounter an error.
- BIOS boot VMs may fail to network boot if a vlan priority tag is set in the DHCP response.

Apr 9, 2026

Checksum: 25c82a53636a5344189fe29c2fc7c950ecf96cb01bc8cccd4cd3af6605dfedae

These updates include the following improvements:

- Update OpenSSL to v3.5.5
- Improves the ability of a pool using GFS2 SRs to fence specific hosts which are having networking issues, rather than allowing them to negatively affect the whole cluster.
- Updates the Secure Boot Forbidden Signature Database (DBX) provided to new VMs to the latest published version.
- Add broadcom-mpi3mr driver.
- When using a VM with disks on a GFS2 or XFS SR, a serial number is now supplied.

These updates include fixes for the following issues:

- When the License Server is unreachable while a XenServer host starts up, the host may not become enabled.
- Certain records (e.g. PCI data) are left behind in the pool database after ejecting a host from a pool.
- A toolstack restart after an update synchronization may block application of updates until another synchronization is performed.
- When installing a UEFI VM from a local DVD, the VM may hang at the firmware screen before starting installation.
- In rare circumstances, the device model can fail due to a problem managing the Xen map cache.
- After upgrading an HA-enabled GFS2 pool from an older XenServer release, HA may not be able to be re-enabled.
- When using an LIO iSCSI target, multipath failover may be delayed.

Mar 18, 2026

Checksum: 0e3647a4f77c8f52e9df18776988a3df99809565d3231cbdc7b6b68a79710a76

These updates include the following improvements:

- Add qllogic-qla2xxx driver

These updates include upstream code changes for the following issues:

- CVE-2024-12243
- CVE-2025-32990
- CVE-2025-32988
- CVE-2025-32989
- CVE-2025-6395
- CVE-2025-9820
- CVE-2025-14831
- CVE-2026-1584

Mar 17, 2026

Checksum: 7ef0e9139be7686bf2b8ab3fe6e2323c5c8e9c1625e94b63dcd bfe82ef17f8c8

These updates include security fixes. For more information, see the [Security Bulletin](#).

These updates include the following improvements:

- Optimizations to VM migration process to reduce migration time for heavily loaded VMs on systems using AMD processors.

Mar 5, 2026

Checksum: fcdc08d17f088b98e6326f8b1173b15803b5135e5759c3f34c3f631157607358

These updates contain foundational changes.

Feb 26, 2026

Checksum: bb6f87ab348e9bb0ce519aac852952f13bf9e59231f67447d41511a77c69dc02

These updates include fixes for the following issues:

- In some situations, a VM export operation can fail to complete.
- A VDI copy operation initiated by a user with the “VM Power Admin” role may incorrectly result in a permissions error.

Fixed issues

February 18, 2026

The following issues have been fixed in XenServer 9. In addition to these fixed issues, additional fixes are released on the [Normal](#) and [Early Access](#) update channels.

XenCenter

For information about known and fixed issues in XenCenter, see [XenCenter What's New](#).

Known Issues

June 26, 2026

This article contains advisories and minor issues in the XenServer 9 release and any workarounds that you can apply.

General

- When attempting to use the serial console to connect to a XenServer® host, the serial console might refuse to accept keyboard input. If you wait until after the console refreshes twice, the console then accepts keyboard input.(CA-311613)

- After a License Activation Service (LAS) disconnection, a XenServer host using online activation might remain in the grace license period even after the LAS becomes reachable again. The host continues to report `LICENSE_SERVER_UNREACHABLE` and `GRACE_LICENSE` alerts despite successfully communicating with the LAS. This issue affects only hosts configured for online activation; hosts using offline activation are unaffected. To work around this issue, restart the toolstack on each affected host. (XSI-2204)
- After applying the June 2026 XenServer updates to hosts originally installed from a XenServer 9 Preview ISO dated May 2026 or earlier, some VM lifecycle operations might fail with an internal error. Affected operations include suspend, checkpoint, import, and storage live migration. This issue occurs because applying updates installs a new kernel but does not reboot the host, so the host continues to run the previous kernel. To work around this issue, reboot the host after applying the update. Hosts installed from June 2026 (or later) ISOs are not affected. (CA-428974)

Guests

- When defining the vCPU topology for a new VM, ensure that the number of cores per socket is set to a power of two (1, 2, 4, 8, and so on). If this value is not set to a power of two, some of your cores are not utilized by the VM, topology information shown inside the VM operating system might be incorrect, and other unexpected behaviors might be observed.(XSI-1766)
- If your VM enters sleep mode, your mouse might become unresponsive although your keyboard remains functional. Rebooting the VM or disabling and enabling the USB controller through Device Manager restores mouse functionality.(CA-399746)
- If a VM is assigned multiple NVIDIA T4 vGPUs, attempts to suspend (or live migrate) the VM might fail.(CA-428435)
- If you assign multiple USB passthrough devices to a VM, the VM might fail to start.(CA-428550)
- On a VM that uses BIOS boot mode, the GRUB boot menu does not respond to keystrokes from the VM console. You cannot use the arrow keys to change the highlighted entry, interrupt the countdown, or select a non-default entry such as a rescue entry. The VM boots the default entry when the countdown ends. This issue affects both newly created and imported BIOS-mode VMs. VMs that use UEFI boot mode are not affected, and the console responds normally after the VM has booted. Where possible, create or import VMs in UEFI boot mode to avoid this issue.(CA-428924)

Performance

- NUMA node drift: when VMs fully use all CPUs on the host, then the hypervisor may run some vCPUs on NUMA nodes far from where the VM's memory is located, causing a performance degra-

dation. Use the [NUMA non-affine vCPUs](#) RRD (XenCenter Performance Graph) to monitor whether you are affected by this, and try to leave at least 1 CPU idle to recover. (XSI-1573)

Windows guests

- For domain-joined Windows 10 VMs (1903 and later) with FireEye Agent installed, repeated successful RDP connections can cause the VM to freeze with 100% CPU usage in [ntoskrnl.exe](#). Perform a hard reboot on the VM to recover from this state. (CA-323760)
- When you create a UEFI VM, the Windows installation requires a key press to start. If you do not press a key during the required period, the VM console switches to the UEFI shell.

To work around this issue, you can restart the installation process in one of the following ways:

```
1 - In the UEFI console, type the following commands:
```

```
1     EFI:
2     EFI\BOOT\BOOTX64
```

```
1 - Reboot the VM
```

When the installation process restarts, watch the VM console for the installation prompt. When the prompt appears, press any key. (CA-333694)

- If you boot a large number of Windows Server 2025 VMs on a single host at the same time, some might fail to start. (CP-51655)
- If an imported Windows VM already has XenServer VM Tools installed, attempting to install the same version of the tools again during the import process fails. (CA-413698)
- If you import a Windows Server VM with more than 4 disks, some of the disks show as offline. (CA-414047)

For details on the latest XenServer VM Tools for Windows release, see [What's new](#).

Linux guests

- You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Red Hat Enterprise Linux 10 (preview), Rocky Linux 8, Rocky Linux 9, CentOS Stream 9, or CentOS Stream 10 VMs as these operating systems do not support memory ballooning with the Xen® hypervisor. (CA-378797)
- On some Linux VMs, especially busy systems with outstanding disk I/O, attempts to suspend or live migrate the VM might fail. To work around this issue, try increasing the value of [/sys/power/pm_freeze_timeout](#), for example, to 300000. If this work-around is not successful, you can upgrade the Linux kernel of the VM to the latest version. (CP-41455)

Storage

- You cannot run more than 200 PVS-Accelerator-enabled VMs on a XenServer host.(CP-39386)
- When read caching is enabled, it may be slower to read from the parent snapshot than from the leaf. To work around this issue, allocate more memory to dom0.(CP-32853)
- Sometimes, when the most recent snapshot for a VM is deleted, it is not possible to completely merge the data from the active VM disk into its parent device. Often, this is caused by the VM being active and writing data to the disk faster than the merge process can merge it. As a result, the VM might use more storage space than is usually necessary due to having two virtual disks in a delta parent/child arrangement. This issue can have more impact on LVM-based SR types as the VMs active write disk is fully provisioned. To allow for the data merge process to complete, either
 - Temporarily reduce the storage I/O rate that the VM is producing.
 - Shut down the VM.

Then rescan the SR on which the VM's disks are stored.(XSI-1922)

Third-party

- A limitation in recent SSH clients means that SSH does not work for user names that contain any of the following characters: { } [] | &. Ensure that your user names and Active Directory server names do not contain any of these characters.

XenCenter

For information about known and fixed issues in XenCenter, see [XenCenter What's New](#).

Deprecation

June 10, 2026

The announcements in this article give you advanced notice of platforms, products, and features that are being phased out so that you can make timely business decisions. We monitor customer use and feedback to determine when they are withdrawn. Announcements can change in subsequent releases and might not include every deprecated feature or functionality.

For details about product lifecycle support, see the [XenServer lifecycle](#) article.

Deprecations and removals

The following table shows the platforms, products, and features that are deprecated or removed.

Deprecated items are not removed immediately. We continue to support them in the current release, but they will be removed in a future release.

Removed items are either removed, or are no longer supported, in XenServer.

Dates in **bold** face indicate changes at this release.

Item	Deprecation announced in	Removed in	Alternative
C support in the XenServer SDK	XenServer 8.4	XenServer 9	
File based licensing	XenServer 8.4		License Activation Service (LAS) is the next generation technology for product activations across the suite of Citrix products. As of April 15, 2026, LAS is the only way to activate and license Citrix on-premises components.
Linux bridge network stack	XenServer 8.4	XenServer 9	vSwitch networking is the only networking technology supported by XenServer.
BIOS boot mode for virtual machines	XenServer 8.4		Install your VMs using UEFI or UEFI Secure Boot mode instead.

Item	Deprecation announced in	Removed in	Alternative
Python 2 scripts in dom0	XenServer 8.4	XenServer 9	Installing additional code or packages or modifying the configuration is not supported unless explicitly instructed by official XenServer documentation or under the guidance of the Support Team. If you are using Python 2 in dom0—for example, in custom power-on scripts—update your code to use Python 3.
BIOS boot mode for XenServer hosts	XenServer 8.4	XenServer 9	Install your hosts with UEFI boot mode instead.
Software Fibre Channel over Ethernet (FCoE)	8.2 CU1	XenServer 9	
Legacy drivers: qla4xxx , qla3xxx , netxen_nic , qlge , qlcnict	8.0	XenServer 9	
FastLinQ drivers: qed , qede , qedf , qedi , qedr	XenServer 8.4	XenServer 9	

Quick start

June 30, 2026

This article steps through how to install and configure XenServer and its graphical, Windows-based

user interface, XenCenter. After installation, it takes you through creating Windows virtual machines (VMs) and then making customized VM templates you can use to create multiple, similar VMs quickly. Finally, this article shows how to create a pool of hosts, which provides the foundation to migrate running VMs between hosts using live migration.

Focusing on the most basic scenarios, this article aims to get you set up quickly.

This article is primarily intended for new users of XenServer and XenCenter. It is intended for those users who want to administer XenServer by using XenCenter. For information on how to administer XenServer using the Linux-based `xe` commands through the XenServer command line interface (known as the `xe` CLI), see [Command-line interface](#).

Terminology and abbreviations

- *Server*: a physical computer that runs XenServer.
- *Host*: a XenServer installation that hosts Virtual Machines (VMs).
- *Virtual Machine (VM)*: a computer composed entirely of software that can run its own operating system and applications as if it were a physical computer. A VM behaves exactly like a physical computer and contains its own virtual (software-based) CPU, RAM, hard disk, and NIC.
- *Pool*: a single managed entity that binds together multiple XenServer hosts and their VMs.
- *Pool coordinator (formerly pool master)*: main host in a pool that provides a single point of contact for all hosts in the pool, routing communication to other members of the pool as necessary.
- *Storage Repository (SR)*: a storage container in which virtual disks are stored.

Major components

XenServer

XenServer is a complete server virtualization platform, with all the capabilities required to create and manage a virtual infrastructure. XenServer is optimized for both Windows and Linux virtual servers.

XenServer runs directly on server hardware without requiring an underlying operating system, which results in an efficient and scalable system. XenServer abstracts elements from the physical machine (such as hard drives, resources, and ports) and allocating them to the virtual machines (VMs) running on it.

XenServer lets you create VMs, take VM disk snapshots, and manage VM workloads.

XenCenter

XenCenter is a graphical, Windows-based user interface. XenCenter enables you to manage XenServer hosts, pools, and shared storage. Use XenCenter to deploy, manage, and monitor VMs from your Windows desktop machine.

The XenCenter *Online Help* is also a great resource for getting started with XenCenter. Press F1 at any time to access context-sensitive information.

Install XenServer and XenCenter

In this section, you set up a minimum XenServer installation.

What you'll learn

You'll learn how to:

- Install XenServer on a single physical server
- Install XenCenter on a Windows computer
- Connecting XenCenter and XenServer to form the infrastructure for creating and running virtual machines (VMs)

Requirements

To get started, you need the following items:

- A physical computer to be the XenServer host
- A Windows computer to run the XenCenter application
- Installation files for XenServer and XenCenter

The XenServer host computer is dedicated entirely to the task of running XenServer and hosting VMs, and is not used for other applications. The computer that runs XenCenter can be any general-purpose Windows computer that satisfies the hardware requirements. You can use this computer to run other applications too. For more information, see [System Requirements](#).

You can download the installation files from [XenServer Downloads](#).

Install the XenServer host

[This is an embedded video. Click the link to watch the video](#)

All servers have at least one IP address associated with them. To configure a static IP address for the host (instead of using DHCP), have the static IP address on hand before beginning this procedure.

Tip:

Press **F12** to advance quickly to the next installer screen. For general help, press **F1**.

To install the XenServer host:

1. Burn the installation files for XenServer to a CD or create a bootable USB.

Note:

For information about using HTTP, FTP, or NFS as your installation source, see [Install XenServer](#).

2. Back up data you want to preserve. Installing XenServer overwrites data on any hard drives that you select to use for the installation.
3. Insert the installation media into the system.
4. Restart the system.
5. Boot from the local installation media (if necessary, see your hardware vendor documentation for information on changing the boot order).
6. Following the initial boot messages and the **Welcome to XenServer** screen, select your keyboard layout for the installation.
7. When the **Welcome to XenServer Setup** screen is displayed, select **Ok**.
8. Read and accept the XenServer EULA.

Note:

If you see a **System Hardware** warning, ensure hardware virtualization assist support is enabled in your system firmware.

9. Select **Ok** to do a clean installation.
10. If you have multiple hard disks, choose a Primary Disk for the installation. Select **Ok**.
Choose which disks you want to use for virtual machine storage. Choose **Ok**.
11. Select **Local media** as your installation source.
12. Select **Skip Verification**, and then choose **Ok**.

Note:

If you encounter problems during installation, verify the installation source.

13. Create and confirm a root password, which the XenCenter application uses to connect to the XenServer host.

14. Set up the management interface to use to connect to XenCenter.

If your computer has multiple NICs, select the NIC which you want to use for management traffic (typically the first NIC).

15. Configure the Management NIC IP address with a static IP address or use DHCP.

16. Specify the host name and the DNS configuration manually or automatically through DHCP.

If you manually configure the DNS, enter the IP addresses of your primary (required), secondary (optional), and tertiary (optional) DNS servers in the fields provided.

17. Select your time zone.

18. Specify how you want the host to determine local time: using NTP or manual time entry. Choose **Ok**.

- If using NTP, you can specify whether DHCP sets the time server. Alternatively, you can enter at least one NTP host name or IP address in the following fields.
- If you selected to set the date and time manually, you are prompted to do so.

19. Select **Install XenServer**.

The installation process starts. This might take some minutes.

20. The next screen asks if you want to install any supplemental packs. Choose **No** to continue.

21. From the **Installation Complete** screen, eject the installation media, and then select **Ok** to reboot the host.

After the host reboots, XenServer displays **xsconsole**, a system configuration console.

Note:

Make note of the IP address displayed. You use this IP address when you connect XenCenter to the host.

Install XenCenter

XenCenter is typically installed on your local system. You can download the XenCenter installer from the [XenServer downloads page](#)

To install XenCenter:

1. Download or transfer the XenCenter installer to the computer that you want to run XenCenter.
2. Double-click the installer `.msi` file to begin the installation.
3. Follow the Setup wizard, which allows you to modify the default destination folder and then to install XenCenter.


Connect XenCenter to the XenServer host

This procedure enables you to add a host to XenCenter.


To connect XenCenter to the XenServer host:

1. Launch XenCenter.

The program opens to the **Home** tab.



2. Click the **Add New Server** icon to open the **Add New Server** dialog box.



✕ Add New Server ? ✕

Enter the host name or IP address of the server you want to add and your user login credentials for that server.

Server:

User login credentials

Username:

Password:

Add Cancel

3. In the **Server** field, enter the IP address of the host. Enter the root user name and password that you set during XenServer installation. Choose **Add**.

Note:

The first time you add a host to XenCenter, the following dialog boxes appear:

- **Save and Restore Connection State** - This dialog box enables you to set your preferences for storing your host connection information and automatically restoring host connections.
- **New Security Certificate** - This dialog box enables you to view the security certificate of the host you are adding to XenCenter. To ensure that you are connecting to the legitimate host, compare the TLS key fingerprint displayed in the certificate details to the TLS key fingerprint obtained directly from the physical host.

License XenServer

You can use XenServer without a license (Trial Edition) for up to 90 days. The Trial Edition lets you try Premium Edition features, but in a restricted size pool of up to 3 hosts. After the 90-day trial period expires, you cannot start any VMs on your Trial Edition host or pool.

If you are using a Citrix License Server that hosts license files to license your hosts and pools, you must change to using licenses held within the License Activation Service in Citrix Cloud. For more information, see [Licensing](#).

For more information about getting a XenServer license, see the [XenServer website](#). XenServer is now an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more](#).

If you have a XenServer license, apply it now. For more information, see [Licensing](#).

Create a pool of XenServer hosts

A resource pool is composed of multiple XenServer host installations, bound together as a single managed entity.

Resource pools enable you to view multiple hosts and their connected shared storage as a single unified resource. You can flexibly deploy VMs across the resource pool based on resource needs and business priorities. A pool can contain up to 32 hosts running the same version of XenServer software, at the same patch level, and with broadly compatible hardware.

One host in the pool is designated as the *pool coordinator*. The pool coordinator provides a single point of contact for the whole pool, routing communication to other members of the pool as necessary. Every member of a resource pool contains all the information necessary to take over the role of pool coordinator if necessary. The pool coordinator is the first host listed for the pool in the XenCenter Resources pane. You can find the pool coordinator's IP address by selecting the pool coordinator and clicking the **Search** tab.

In a pool with shared storage, you can start VMs on *any* pool member that has sufficient memory and dynamically move the VMs between hosts. The VMs are moved while running and with minimal downtime. If an individual XenServer host suffers a hardware failure, you can restart the failed VMs on another host in the same pool.

If the high availability feature is enabled, protected VMs are *automatically* moved if a host fails. On an HA-enabled pool, a new pool coordinator is automatically nominated if the pool coordinator is shut down.

Note:

For a description of heterogeneous pool technology, see [Hosts and resource pools](#).

What you'll learn

You'll learn how to:

- Create a pool of hosts
- Set up a network for the pool
- Bond NICs
- Set up shared storage for the pool

While XenServer accommodates many shared storage solutions, this section focuses on two common types: NFS and iSCSI.

Requirements

To create a pool with shared storage, you need the following items:

- A second XenServer host, with similar processor type.
Connect this host to your XenCenter application.
- A storage repository for IP-based storage
- A consistent network configuration for all hosts you want to add to your pool. All hosts must have the same number of NICs. NIC0 on each host in the pool must be on the same network.

To get you started quickly, this section focuses on creating *homogeneous* pools. Within a homogeneous pool, all hosts must have compatible processors and be running the same version of XenServer, under the same type of XenServer product license. For a full list of homogeneous pool requirements, see [System requirements](#).

Create a pool


To create a pool:

1. On the toolbar, click the **New Pool** button.



2. Enter a name and optional description for the new pool.
3. Nominate the pool coordinator by selecting a host from the **Coordinator** list.
4. Select the second host to place in the new pool from the **Additional members** list.
5. Click **Create Pool**.

The new pool appears in the **Resources** pane.



Set up networks for the pool

When you install XenServer, you create a network connection, typically on the first NIC in the pool where you specified an IP address (during XenServer installation).

However, you may need to connect your pool to VLANs and other physical networks. To do so, you must add these networks to the pool. You can configure XenServer to connect each NIC to one physical network and numerous VLANs.

For detailed information about networking configuration, see [Create networks in resource pools](#).

Before creating networks, ensure that the cabling matches on each host in the pool. Plug the NICs on each host into the same physical networks as the corresponding NICs on the other pool members.

Note:


If the NICs were not plugged in to the NICs on the host when you installed XenServer:

- Plug the NICs in
- In XenCenter, select **<your host> > NICs** tab
- Click **Rescan** for them to appear

For additional information about configuring XenServer networking, see [Networking](#) and [About XenServer Networks](#).

To add a network to XenServer:

1. In the **Resources** pane in XenCenter, select the pool.
2. Click the **Networking** tab.
3. Click **Add Network**.



4. On the **Select Type** page, select **External Network**, and click **Next**.
5. On the **Name** page, enter a meaningful name for the network and description.
6. On the **Network settings** page, specify the following:
 - **NIC:** Select the NIC that you want XenServer to use to send and receive data from the network.
 - **VLAN:** If the network is a VLAN, enter the VLAN ID (or “tag”).
 - **MTU:** If the network uses jumbo frames, enter a value for the Maximum Transmission Unit (MTU) between 1500 to 9216. Otherwise, leave the MTU box at its default value of 1500.

If you configure many virtual machines to use this network, you can select the **Automatically add this network to new virtual machines** check box. This option adds the network by default.

7. Click **Finish**.

Bonding NICs

NIC bonding can make your host more resilient by using two or more physical NICs as if they were a single, high-performing channel. This section only provides a very brief overview of bonding, also known as *NIC teaming*. Before configuring bonds for use in a production environment, we recommend reading more in-depth information about bonding. For more information, see [Networking](#).

XenServer supports the following bond modes: Active/active, active/passive (active/backup), and LACP. Active/active provides load balancing and redundancy for VM-based traffic. For other types of traffic (storage and management), active/active cannot load balance traffic. As a result, LACP or multipathing are better choice for storage traffic. For information about multipathing, see [Storage](#). For more information about bonding, see [Networking](#).

LACP requires that your switches support the IEEE 802.3ad standard. The switch must contain a separate LAG group configured for each LACP bond on the host. For more details about creating LAG groups, see [Networking](#).

To bond NICs:

1. Ensure that the NICs you want to bind together are not in use: shut down any VMs with virtual network interfaces using these NICs before creating the bond. After you have created the bond, reconnect the virtual network interfaces to an appropriate network.
2. Select the host in the **Resources** pane then open the **NICs** tab and click **Create Bond**.
3. Select the NICs you want to bond together. To select a NIC, select its check box in the list. Up to four NICs may be selected in this list. Clear the check box to deselect a NIC. To maintain a flexible and secure network, you can bond either two, three, or four NICs.

4. Under **Bond** mode, choose the type of bond:

- Select **Active-active** to configure an active-active bond. Traffic is balanced between the bonded NICs. If one NIC within the bond fails, the host's network traffic automatically routes over the second NIC.
- Select **Active-passive** to configure an active-passive bond. Traffic passes over only one of the bonded NICs. In this mode, the second NIC only becomes active if the active NIC fails, for example, if it loses network connectivity.
- Select **LACP with load balancing based on source MAC address** to configure a LACP bond. The outgoing NIC is selected based on MAC address of the VM from which the traffic originated. Use this option to balance traffic in an environment where you have several VMs on the same host. This option is not suitable if there are fewer virtual interfaces (VIFs) than NICs: as load balancing is not optimal because the traffic cannot be split across NICs.
- Select **LACP with load balancing based on IP and port of source and destination** to configure a LACP bond. The source IP address, source port number, destination IP address, and destination port number are used to allocate the traffic across the NICs. Use this option to balance traffic from VMs in an environment where the number of NICs exceeds the number of VIFs.

5. To use jumbo frames, set the Maximum Transmission Unit (MTU) to a value between 1500 to 9216.

6. To have the new bonded network automatically added to any new VMs created using the New VM wizard, select the check box.

7. Click **Create** to create the NIC bond and close the dialog box.

XenCenter automatically moves management and secondary interfaces from secondary bonded NICs to the bond interface when the new bond is created. A host with its management interface on a bond is not permitted to join a pool. Before the host can join a pool, you must reconfigure the management interface and move it back on to a physical NIC.

Setting up shared storage for the pool

To connect the hosts in a pool to a remote storage array, create a XenServer SR. The SR is the storage container where a VM's virtual disks are stored. SRs are persistent, on-disk objects that exist independently of XenServer. SRs can exist on different types of physical storage devices, both internal and external. These types include local disk devices and shared network storage.

You can configure a XenServer SR for various different types of storage, including:

- NFS

- Software iSCSI
- Hardware HBA
- SMB
- Fibre Channel
- Software FCoE (deprecated)

This section steps through setting up two types of shared SRs for a pool of hosts: NFS and iSCSI. Before you create an SR, configure your NFS or iSCSI storage array. Setup differs depending on the type of storage solution that you use. For more information, see your vendor documentation. Generally, before you begin, complete the following setup for your storage solution:

- **iSCSI SR:** You must have created a volume and a LUN on the storage array.
- **NFS SR:** You must have created the volume on the storage device.
- **Hardware HBA:** You must have done the configuration required to expose the LUN before running the **New Storage Repository** wizard
- **Software FCoE SR:** You must have manually completed the configuration required to expose a LUN to the host. This setup includes configuring the FCoE fabric and allocating LUNs to your SAN's public world wide name (PWWN).

If you are creating an SR for IP-based storage (iSCSI or NFS), you can configure one of the following as the storage network: the NIC that handles the management traffic or a new NIC for the storage traffic. To configure a different a NIC for storage traffic, assign an IP address to a NIC by creating a *management interface*.

When you create a management interface, you must assign it an IP address that meets the following criteria:

- The IP address is on the same subnet as the storage controller, if applicable
- The IP address is on a different subnet than the IP address you specified when you installed XenServer
- The IP address is not on the same subnet as any other management interfaces.

To assign an IP address to a NIC:

1. Ensure that the NIC is on a separate subnet or that routing is configured to suit your network topology. This configuration forces the desired traffic over the selected NIC.
2. In the **Resource** pane of XenCenter, select the pool (or standalone host). Click the **Networking** tab, and then click the **Configure** button.
3. In the **Configure IP Address** dialog, in the left pane, click **Add IP address**.

4. Give the new interface a meaningful name (for example, *yourstoragearray_network*). Select the **Network** associated with the NIC that you use for storage traffic.
5. Click **Use these network settings**. Enter a static IP address that you want to configure on the NIC, the subnet mask, and gateway. Click **OK**. The IP address must be on the same subnet as the storage controller the NIC is connected to.

Note:


Whenever you assign a NIC an IP address, it must be on a different subnet than any other NICs with IP addresses in the pool. This includes the primary management interface.

To create a new shared NFS or iSCSI storage repository:

1. On the **Resources** pane, select the pool. On the toolbar, click the **New Storage** button.



The **New Storage Repository** wizard opens.



2. Under Virtual disk storage, choose NFS or iSCSI as the storage type. Click Next to continue.
3. If you choose NFS:
 - a) Enter a name for the new SR and the name of the share where it is located. Click **Scan** to have the wizard scan for existing NFS SRs in the specified location.

Note:

The NFS host must be configured to export the specified path to all XenServer hosts in the pool.

- b) Click **Finish**.

The new SR appears in the **Resources** pane, within the pool.

4. If you choose iSCSI:

- a) Enter a name for the new SR and then the IP address or DNS name of the iSCSI target.

Note:

The iSCSI storage target must be configured to enable every XenServer host in the pool to have access to one or more LUNs.

- b) If you have configured the iSCSI target to use CHAP authentication, enter the user name and password.
- c) Click the **Scan Target Host** button, and then choose the iSCSI target IQN from the Target IQN list.

Warning:

The iSCSI target and all hosts in the pool must have *unique* IQNs.

- d) Click **Target LUN**, and then select the LUN on which to create the SR from the Target LUN list.

Warning:

Each individual iSCSI storage repository must be contained entirely on a single LUN and cannot span more than one LUN. Any data present on the chosen LUN is destroyed.

- e) Click **Finish**.

The new SR appears in the **Resources** pane, within the pool.

The new shared SR now becomes the default SR for the pool.

Create virtual machines

Through XenCenter, you can create virtual machines in various ways, according to your needs. Whether you are deploying individual VMs with distinct configurations or groups of multiple, similar VMs, XenCenter gets you up and running in just a few steps.

XenServer also provides an easy way to convert batches of virtual machines from VMware. For more information, see [Conversion Manager](#).

This section focuses on a few methods by which to create Windows VMs. To get started quickly, the procedures use the simplest setup of XenServer: a single XenServer host with local storage (after you connect XenCenter to the XenServer host, storage is automatically configured on the local disk of the host).

This section also demonstrates how to use live migration to live migrate VMs between hosts in the pool.

After explaining how to create and customize your new VM, this section demonstrates how to convert that existing VM into a VM template. A VM template preserves your customization so you can always use it to create VMs to the same (or to similar) specifications. It also reduces the time taken to create multiple VMs.

You can also create a VM template from a snapshot of an existing VM. A snapshot is a record of a running VM at a point in time. It saves the storage, configuration, and networking information of the original VM, which makes it useful for backup purposes. Snapshots provide a fast way to make VM templates. This section demonstrates how to take a snapshot of an existing VM and then how to convert that snapshot into a VM template. Finally, this section describes how to create VMs from a VM template.

What you'll learn

You'll learn how to:

- Create a Windows 10 VM
- Install XenServer VM Tools for Windows
- Migrate a running VM between hosts in the pool
- Create a VM template
- Create a VM from a VM template

Requirements

To create a pool with shared storage, you need the following items:

- The XenServer pool you set up
- XenCenter
- Installation files for Windows 10 VM
- Installation files for XenServer VM Tools for Windows

Create a Windows 10 VM

Note:

The following procedure provides an example of creating Windows 10 VM. The default values may vary depending on the operating system that you choose.

To create a Windows VM:

1. On the toolbar, click the **New VM** button to open the **New VM** wizard.



The **New VM** wizard allows you to configure the new VM, adjusting various parameters for CPU, storage, and networking resources.

2. Select a VM template and click **Next**.

Each template contains the setup information for creating a VM with a specific guest operating system (OS), and with optimum storage. This list reflects the templates that XenServer currently supports.

Note:

If the OS you're installing on your new VM is compatible only with the original hardware, check the **Copy host BIOS strings to VM** box. For example, use this option for an OS installation CD that was packaged with a specific computer.

After you first start a VM, you cannot change its BIOS strings. Ensure that the BIOS strings are correct before starting the VM for the first time.

3. Enter a name for and optional description of the new VM.
4. Choose the source of the OS media to install on the new VM.

Installing from a CD/DVD is the simplest option for getting started. Choose the default installation source option (DVD drive), insert the disk into the DVD drive of the XenServer host, and choose **Next** to proceed.

XenServer also allows you to pull OS installation media from a range of sources, including a pre-existing ISO library.

To attach a pre-existing ISO library, click **New ISO library** and indicate the location and type of ISO library. You can then choose the specific operating system ISO media from the list.

5. The VM runs on the installed host. Choose **Next** to proceed.
6. Allocate processor and memory resources.

Each OS has different configuration requirements which are reflected in the templates. You can choose to modify the defaults if necessary. Click **Next** to continue.

7. Assign a graphics processing unit (GPU).

The **New VM** wizard prompts you to assign a dedicated GPU or virtual GPUs to the VM. This option enables the VM to use the processing power of the GPU. It provides better support for high-end 3D professional graphics applications such as CAD, GIS, and Medical Imaging applications.

Note:

GPU Virtualization is available for XenServer Premium Edition customers.

8. Configure storage for the new VM.

Click **Next** to select the default allocation and configuration, or you might want to:

- a) Change the name, description, or size of your virtual disk by clicking **Edit**.
- b) Add a new virtual disk by selecting **Add**.

Note:

When you create a pool of XenServer hosts, you can configure shared storage at this point when creating a VM.

9. Configure networking on the new VM.


Click **Next** to select the default NIC and configurations, including an automatically created unique MAC address for each NIC, or you can:

- a) Change the physical network, MAC address, or Quality of Service (QoS) priority of the virtual disk by clicking **Edit**.
- b) Add a new virtual network interface by selecting **Add**.

XenServer uses the virtual network interface to connect to the physical network on the host. Be sure to select the network that corresponds with the network the virtual machine requires. To add a physical network, see [Setting Up Networks for the Pool](#).

10. Review settings, and then click **Create Now** to create the VM and return to the **Search** tab.

An icon for your new VM appears under the host in the **Resources** pane.



On the **Resources** pane, select the VM, and then click the **Console** tab to see the VM console.

11. Follow the OS installation screens and make your selections.
12. After the OS installation completes and the VM reboots, install the XenServer VM Tools for Windows.

Install XenServer VM Tools for Windows

XenServer VM Tools for Windows provide high performance I/O services without the overhead of traditional device emulation. XenServer VM Tools for Windows consists of I/O drivers (also known as paravirtualized drivers or PV drivers) and the Management Agent. XenServer VM Tools for Windows must be installed on each Windows VM for the VM to have a fully supported configuration. A Windows VM functions without them, but performance is hampered. XenServer VM Tools for Windows also enable certain functions and features, including cleanly shutting down, rebooting, suspending and live migrating VMs.

Warning:

Install XenServer VM Tools for Windows for each Windows VM. Running Windows VMs without XenServer VM Tools for Windows is *not* supported.

We recommend that you snapshot your VM before installing or updating the XenServer VM Tools.

To install XenServer VM Tools for Windows:

1. Download the XenServer VM Tools for Windows file onto your Windows VM. Get this file from the [XenServer downloads page](#).
2. Run the `managementagentx64.msi` file to begin XenServer VM Tools installation.
3. Follow the prompts in the installer.
4. Restart the VM when prompted to complete the installation process.

Note:

I/O drivers are automatically installed on a Windows VM that can receive updates from Windows Update. However, we recommend that you install the XenServer VM Tools for Windows package to install the Management Agent, and to maintain a supported configuration. The following features are available only for XenServer Premium Edition customers:

- Ability to receive I/O drivers from Windows Update
- Automatic updating of the Management Agent

After you have installed the XenServer VM Tools for Windows, you can customize your VM by installing applications and performing any other configurations. If you want to create multiple VMs with similar specifications, you can do so quickly by making a template from the existing VM. Use that template to create VMs. For more information, see [Creating VM Templates](#).

Migrate running VMs between hosts in a pool

Using live migration, you can move a running VM from one host to another in the same pool, and with virtually no service interruption. Where you decide to migrate a VM to depends on how you configure the VM and pool.

To migrate a running VM:

1. On the **Resources** pane, select the VM that you want to move.


Note:

Ensure that the VM you migrate does not have local storage.

2. Right-click the VM icon, point to **Migrate to Server**, and then select the new VM host.

Tip:

You can also drag the VM onto the target host.



3. The migrated VM displays under the new host in the **Resources** pane.

Create VM templates

There are various ways to create a VM template from an existing Windows VM, each with its individual benefits. This section focuses on two methods: converting an existing VM into a template, and creating a template from a snapshot of a VM. In both cases, the VM template preserves the customized configuration of the original VM or VM snapshot. The template can then be used to create new, similar VMs quickly. This section demonstrates how to make new VMs from these templates.

Before you create a template from an existing VM or VM snapshot, we recommend that you run the Windows utility **Sysprep** on the original VM. In general, running **Sysprep** prepares an operating system for disk cloning and restoration. Windows OS installations include many unique elements per installation (including Security Identifiers and computer names). These elements must stay unique and not be copied to new VMs. If copied, confusion and problems are likely to arise. Running **Sysprep** avoids these problems by allowing the generation of new, unique elements for the new VMs.

Note:

Running **Sysprep** may not be as necessary for basic deployments or test environments as it is for production environments.

For more information about **Sysprep**, see your Windows documentation. The detailed procedure of

running this utility can differ depending on the version of Windows installed.

Create a VM template from an existing VM To create a VM template from an existing VM:

Warning:


When you create a template from an existing VM, the new template replaces the original VM. The VM no longer exists.

1. Shut down the VM that you want to convert.
2. On the **Resources** pane, right-click the VM, and select **Convert to Template**.
3. Click **Convert** to confirm.

Once you create the template, the new VM template appears in the **Resources** pane, replacing the existing VM.

Create a VM template from a VM snapshot To create a template from a snapshot of a VM:

1. On the **Resources** pane, select the VM. Click the **Snapshots** tab, and then **Take Snapshot**.
2. Enter a name and an optional description of the new snapshot. Click **Take Snapshot**.
3. Once the snapshot finishes and the icon displays in the **Snapshots** tab, select the icon of the new snapshot. From the **Actions** list, choose **New Template from Snapshot**.



4. Enter a name for the template, and then click **Create**.

Create VMs from a VM template

To create a VM from a customized VM template:

1. On the XenCenter **Resources** pane, right-click the template, and select **New VM** wizard.

The **New VM** wizard opens.

2. Follow the **New VM** wizard to create a VM from the selected template.

Note:

When the wizard prompts you for an OS installation media source, select the default and continue.

The new VM appears in the **Resources** pane.

If you are using a template created from an existing VM, you can also choose to select **Quick Create**. This option does not take you through the **New VM** wizard. Instead this option instantly creates and provisions a new VM using all the configuration settings specified in your template.

XenServer® Enterprise Reference Architecture

May 12, 2026

Overview

This reference architecture defines a recommended approach for designing, deploying, and operating XenServer to support enterprise-scale workloads. It provides a validated foundation for hosting Citrix Virtual Apps and Desktops™ (CVAD) environments, as well as general server virtualization, with a focus on scalability, resilience, and operational simplicity.

The architecture is built around the concept of the XenServer resource pool as the core unit of scale and management. Each resource pool is designed to operate within a single data center (or a set of closely connected data centers) and supports predictable, high-performance workload execution. Deployments can be scaled horizontally by adding additional resource pools to meet growing capacity and organizational requirements.

Design Principles

The following principles summarize the design approach:

- **Resource pool as the unit of scale:** Deploy and manage workloads within well-defined resource pools, and scale horizontally by adding additional pools
- **Workload isolation by design:** Align resource pools to specific workload types to ensure consistent performance and operational behavior
- **N+1 capacity model:** Maintain sufficient capacity to tolerate a single host failure without impacting workload availability
- **Separation of concerns:** Clearly separate management, VM, and storage traffic to ensure performance, resilience, and security
- **Explicit resilience strategy:** Treat intra-data center resilience and inter-data center disaster recovery as distinct architectural concerns
- **Operational simplicity:** Introduce optional capabilities only where they provide clear value, avoiding unnecessary complexity

Together, these principles provide a consistent framework for designing XenServer environments that are scalable, resilient, and straightforward to operate, while remaining adaptable to evolving workload and organizational requirements.

This architecture assumes the use of remote block storage, integrated identity management through Active Directory, and secure management via TLS. Optional capabilities such as Workload Balancing (WLB), High Availability (HA), and Disaster Recovery (DR) can be incorporated based on workload requirements, but each introduces additional operational considerations that must be explicitly planned.

The reference architecture is intended for enterprise environments requiring predictable performance and strong operational control. It is not optimized for specialized scenarios such as very large individual VM disks (>2 TB) or GPU-intensive workloads, which may require alternative designs or adaptations.

Overall, this document serves as both a design guide and an operational framework, enabling organizations to deploy XenServer in a consistent, supportable manner while retaining flexibility to extend and evolve the platform over time.

Assumptions and Scope

The architecture is based on a set of foundational assumptions about scale, infrastructure, and operational practices.

- All hardware used must be listed in the XenServer [Hardware Compatibility List \(HCL\)](#)

- Each resource pool is expected to support up to 1,000 virtual machines, with overall deployments scaling by adding additional pools rather than expanding a single pool indefinitely, up to a maximum of 200 pools. All hosts within a pool are assumed to operate within a tightly coupled network boundary, ensuring consistent and reliable access to shared storage
 - Hosts must have sufficient RAM to support the required workload and Control Domain with any performance caching requirements (see Sizing the Resource Pool), up to a maximum of 6 TB
 - Hosts must either have local storage for XenServer OS (minimum 46 GB, ideally 70 GB), or ability to boot from SAN
 - Hosts must use UEFI firmware; legacy BIOS is not supported in XenServer 9
- Storage is assumed to be provided via remote block-based systems, with XenServer using LVM-based storage repositories. Efficiency and resilience are delivered primarily by the underlying storage platform, including capabilities such as thin provisioning to enable active allocation on use, and multipathing. Continuous and reliable storage connectivity is a critical requirement.
- Networking follows a model of strict separation between management, virtual machine, and storage traffic. To achieve this hosts must be equipped with a minimum of either 2 NICs and 2 Fiber Channel connections, or 4 NICs. This separation is combined with bonded network interfaces to provide resilience and throughput, while avoiding configurations that introduce unnecessary dependencies or performance constraints.
- Security and identity are treated as integral to the design. Management communications are secured using TLS, and integration with Active Directory is assumed for authentication and role-based access control. Administrative access is expected to follow standard enterprise security practices, with restricted use of local accounts and controlled exposure of management interfaces.
- Operationally, resource pools are designed around an N+1 capacity model, allowing workloads to continue running during host maintenance or in the event of a single host failure. Updates and upgrades are expected to be applied regularly to maintain supportability and security. Where disaster recovery is required, it is implemented as a coordinated operational process rather than a seamless or fully automated failover capability.

Optional capabilities can be incorporated where required, but are not assumed by default. These include:

- Workload Balancing (WLB)
- High Availability (HA) for workloads
- Disaster Recovery (DR) across data centers

Each of these introduces additional operational complexity and should be adopted based on clear workload and business requirements.

Out of Scope

This reference architecture is not designed for all scenarios and should not be treated as a universal solution. In particular, it does not directly address:

- Workloads requiring individual virtual disks greater than 2 TB
- GPU-accelerated or GPU-dependent workloads
- Architectures requiring seamless, active-active workload mobility across geographically dispersed data centers

In these cases, elements of this design may still be applicable, but additional architectural considerations and adaptations will be required.

Definitions

Terms and definitions to support the reading and understanding of this reference architecture.

Term	Definition
Data Center	A collection of networked compute which is located closely together and has access to remote storage. All connectivity is expected to be low-latency, high-bandwidth, and highly reliable. In particular, XenServer hosts must not lose connectivity to storage. The compute can be configured to be resilient to failure within the data center and can form part of a disaster recovery solution for other data centers but is expected to operate as a standalone deployment during normal operation. This means there is no expectation of regular migration of workloads outside of the data center to form a wider resilience footprint. Network interconnect between any resources deployed within a data center is expected to have <2 ms latency and >=10 Gbps network throughput.

Term	Definition
Close Proximity Data Centers	Multiple data centers with low-latency, high-bandwidth, and highly reliable, interconnect, expected to operate as if they are a single logical data center. Network interconnect between any resources deployed across these data centers is expected to have <5 ms latency and >=10 Gbps network throughput.
Geo Dispersed Data Centers	Data centers with large distances between them. Such data centers are expected to operate as independent entities, each with their own networks and storage.
Upgrade	A major product version change (for example, upgrading from XenServer 8.4 to XenServer 9)
Update	Installation of packages within a specific single version of XenServer, providing additional features, bug fixes, and security fixes.
Resource Pool (Often shortened to Pool)	A management unit that groups together hosts and provides a single point to manage Storage and Networks across the set of hosts. For more detail see Resource Pools

Architecture Blueprint

The expectations for resource pools are as follows:

- Each resource pool is designed for up to 1000 running VMs.
- Each resource pool is entirely located within a single data center or across close proximity data centers.
- Each resource pool provides for one specific use case, where all the workloads have similar characteristics for availability and performance. For small deployments, mixed use cases are possible but care must be taken when considering the load balancing and failure modes for the resource pools.

Examples of single use-cases are:

- Virtual desktops for a CVAD deployment
- Application Servers for CVAD deployment
- CVAD infrastructure

- General server virtualization workloads

Each resource pool will have the following attributes

- LVM protocol used for all remote block storage. It is not recommended to use GFS2 protocol.
- High availability for Pool Coordinator management
- Use of TLS 1.2 for all management communications
- **Optional:** Workload Balancing (WLB) for workload VMs
- **Optional:** High Availability (HA) of workload VMs
- **Optional:** Disaster Recovery (DR)
- **Optional:** XenServer Conversion Manager installed

Note:

The design can be scaled by creating multiple resource pools within data centers or across data centers to provide the workload where it is needed.

Construct each resource pool as defined in the sections below.

Core Resource Pool configuration

When building out a new resource pool the best practice is to configure a standalone host with network and storage configuration as required, then add other hosts to this pool. The additional hosts will gain the configuration as they are added to the resource pool.

Configure role-based access control Connect all resource pools to a trusted Active Directory domain to allow AD users and groups to be easily managed and audited and to control access and permissions to the XenServer resource pools.

Note:

Do not use the default (root) account for management except in emergency conditions where it is not possible to authenticate via AD credentials (e.g. AD Domain Controller failure).

Provisioning Certificates for Hosts All hosts in all resource pools should have a provisioned TLS certificate for use on the management network. For CVAD workloads the hosts will need to be resolvable by their FQDN in DNS.

All hosts must have a certificate provisioned to them; this may be a shared wildcard certificate or a per host certificate.

Note:

Some applications that use the XenServer APIs, such as some older versions of Citrix Apps and Desktops, require the use of an individual certificate for each host and don't allow the use of wildcard certificates. See the relevant documentation for more details:

- [DaaS](#)
- [CVAD 2507 and later](#)
- [CVAD earlier than 2507](#)

Any 3rd-party connections to the resource pool should use TLS and the FQDN to address the hosts in the resource pool.

Securing the communication to the resource pool

- Disable port 80
- Set SSH to **Auto** mode so that if the management toolstack fails, SSH is automatically enabled

Enable Secure Boot XenServer 9 hosts can be configured to run with Secure Boot enabled in firmware. This enforces signature verification at each stage of the boot process, reducing the risk of untrusted or malicious code executing before the host operating system is fully loaded.

Enable Secure Boot in firmware settings on each host. All Dom0 kernel modules in XenServer 9 are signed, so standard host operation is not affected. Mixed pools where Secure Boot is enabled on some hosts and disabled on others are supported.

For full configuration steps, see [Secure Boot for XenServer 9](#).

Resource Pool Network Configuration

Separation is required for the following networks:

- Management network
- VM network
- Storage network (if using network based storage)

This separation can be provided by separate NICs or by using VLANs using the same NICs. If using Network based storage it is strongly recommended to use separate NICs for performance reasons.

Management and VM networks The Management and VM networks should be operated on bonded NICs to provide resilience and throughput.

The bonds should be configured as LACP bonds where possible to ensure the best response to connectivity loss while maximizing the bandwidth availability. This will require the switches to be configured in a matching configuration. If the switches cannot be configured to support LACP then active-active is beneficial for VM networks. Management networks can use active-passive without any impact. For more information see [Bonding Types](#).

Where possible, the switches should be cross-linked to the NICs to ensure that there is no single point of failure in the network transport and they should use separate power sources or redundant supplies, see [Networking Best Practices](#) for full details.

All the hosts in the pool must have fixed IP addresses on the management network; XenServer does not support the scenario where the host IP address issued by DHCP changes. If fixed DHCP addresses are used, the DHCP environment should be expected to be reliably available (e.g. as part of the network infrastructure) and should not be provided by a VM running on a host that depends on that DHCP server. Static addressing for infrastructure is recommended.

All host management networks need to allow connection to synchronized NTP servers to ensure their clocks stay aligned; this should be the same time source as used for the Active Directory providing the computer and user accounts for the role-based access control to ensure that clock drift does not occur and break the communication paths. This will also ensure that there are consistent timestamps for all system logs.

The full connectivity requirements that are required to enable operation of the environment are listed [here](#).

Storage Networks When using network-based storage the best practice is to use multipath support from the storage, which requires independent networks (i.e. different subnets). This typically means using single NICs rather than bonds and configuring appropriate subnets on each. Consult your storage provider for further recommendations.

Resource Pool Storage Configuration

Multipathing should be enabled. This must be set on all hosts individually, and there should be at least 2 paths available for all storage resources. These paths should be diverse across network/fiber channel infrastructure.

LVM is thick provisioned storage, so XenServer will require the full amount of storage needed for the VM disks created. This means that every disk and template will take the full capacity available even if only small portions of those disks are in use.

The recommendation is to use SANs that provide their own thin provisioning function to ensure efficient use of the storage. The SAN thin provisioning better utilizes the available storage by allocating disk storage space to VDIs as data is written to the virtual disk, rather than allocating the full virtual

size of the VDI in advance. The amount of storage used should be monitored carefully by the administrator using tools provided by the SAN vendor.

If thin provisioning is used on the SAN, monitoring storage utilization is important; if the SAN capacity is exhausted such that further writes are not possible this can lead to VM failures and/or data corruption.

If using MCS to provision workloads for Citrix® environments it is recommended to use a small number of large LUNs for this as the golden image needs to be copied onto all LUNs which can be slow and make image provisioning take significant elapsed time, delaying the roll out of updated images.

Resilience Strategy

XenServer resource pools, along with their storage resources, are operated within a single data center (or within a collection of close proximity data centers).

These resource pools can be configured to be resilient to local server, storage, and network failures within the data center. In addition, resource pools can be used in combination to provide disaster recovery functions in case a data center fails and causes loss of an entire resource pool.

These are considered separately below.

Resilience within a Data Center

While workloads within a data center can be protected through a full DR configuration as described below, there are mechanisms for resilience within each resource pool to ensure high availability of the resources. These mechanisms are designed to be automated and used more frequently to protect against more common and localized issues. The following facilities are available:

- Management High Availability (through automated coordinator elections during unplanned outages)
- Workload compute High Availability (through automated VM restart on unplanned outages)
- Storage resilience (through storage multipathing)
- Network resilience (through network bonding and switch configurations)

To achieve this there is an expectation that the resource pool operates with a minimum of N+1 hosts, where N is the number of hosts required to operate the workload. This provides the capability to operate the resource pool at full capacity with one host unavailable.

Note:

N+1 Host capacity enables updating and upgrading resource pools without workload outage, allowing easy application of security updates and feature enhancements.

These facilities will provide a highly resilient deployment that is operated completely within a data center or across close proximity data centers. It will not support resilience across geographically dispersed data centers. The only way to support this is through the DR function described below (see Resilience across Data Centers).

Optional Capabilities

Workload Balancing (WLB)

- **Use when:** VM workloads have materially different and time-varying utilization patterns during their runtime and you want automated recommendations and/or rebalancing to maintain performance.
- **Avoid when:** Workloads are highly uniform (for example, many similar non-persistent VDI VMs) where initial placement and normal lifecycle operations already provide adequate balance, or where migrations are operationally undesirable.
- **Operational impact:** Adds an additional appliance to operate and monitor; includes a daily maintenance window during which WLB processing is unavailable.

WLB can be used to maintain the best host utilization and performance as workload usage changes over time.

Import the appliance into the resource pool and use the Management network for its connectivity. The Workload Balancing (WLB) appliance may be downloaded from the [XenServer downloads page](#).

The default sizing for the appliance (2 GB Ram, 30 GB disk and 2 vCPUs) should be used.

Workload suitability WLB should only be required when workload demand changes dynamically within VMs over their running lifetime. If WLB is used, every time a VM is stopped and started again it is placed optimally for the current workload immediately, so VMs should only need moving as host utilization changes. When the use case provides many similar VMs all carrying out similar tasks there is little value in using WLB and you can rely on the initial VM placement to provide workload balancing.

Adjust database maintenance window The WLB appliance performs routine maintenance daily. By default, this is at 00:05 UTC. This should be configured to happen when there is low activity in the resource pool as there will be a WLB processing outage for up to 30 mins during this process. This does not affect normal operation, only placement optimization decisions; WLB will complete actions as soon as its maintenance is completed.

WLB parameter configuration Many configuration settings can be left as defaults. This will set up the WLB feature to make recommendations to an administrator, and the administrator will have to apply these actions.

The default and expected configuration for this blueprint only requires the ‘Critical Thresholds’ and weightings to be adjusted to enable the resource pools to operate at maximum performance. These can be optimized through use, but some suggested initial values are as follows:

- **CPU Utilization:** The default for this is 90% critical threshold with a maximum weighting. This will ensure that once a host CPU reaches an average loading of 90% over the last 1.5 minutes, the WLB service will evaluate alternate locations for VMs to be run.
 - If the performance of your VMs is being impacted at lower CPU utilizations, you can adjust this critical threshold with the aim of keeping the average CPU usage at a lower value across hosts where possible.
 - Recommended initial critical threshold Value: **90%**
 - Recommended initial metric weighting value: **Most important**
- **Free Memory:** There is no specific performance to be gained by keeping free memory on hosts, so migrating machines to keep memory free is an overhead that is not required.
 - Recommended initial critical threshold Value: **0**
 - Recommended initial metric weighting value: **Least important**
- **Network Read and write:** The default values for these are 25 MB/s. This means that if any host is utilizing more bandwidth for reads or writes it will be considered for re-location. This parameter is not relevant unless you have some VMs running heavy network workloads, and avoiding impact on other VMs is relevant.
 - For many applications this is not a relevant factor, and the weighting should be set to the least important setting.
 - For many modern networks this is a very low value and if this setting is relevant, we recommend setting this to allow good use of this bandwidth before considering migration to another host.
 - Recommended initial critical threshold Value: **90% of available bandwidth**
 - Recommended initial metric weighting value: **Medium importance**
- **Disk Read and write:** The default values for this are 25 MB/s. This means that if any host is utilizing more bandwidth for reads or writes it will be considered for re-location. This parameter is not relevant unless you have some VMs running heavy disk workloads and avoiding impact on other VMs is relevant. If remote storage is used, then unless the bandwidth of the fabric is the bottleneck this is of no value as any impact on remote storage will impact all VM performance that uses that storage regardless of the host it is being operated on.
 - For most modern storage and fabrics this is a low value and if this setting is relevant, it should be set to a higher value to allow for good use of bandwidth. This should be set high enough that regular out of hours maintenance on VMs does not cause mass migration to occur when not needed.

- Recommended initial critical threshold Value: **90% of available bandwidth**
- Recommended initial metric weighting value: **Medium importance**

High Availability (HA) for the Resource Pool

Full HA documentation is available [here](#)

There are 2 elements to consider for XenServer regarding High Availability

- Management HA: The ability to automatically maintain management operations when there is an unexpected outage of the coordinator host in the resource pool
- VM HA: The ability to maintain workload operations when there is an unexpected outage of a host in the resource pool

Both these aspects require specific storage and Network configurations. This is covered in the sections below.

One of the Storage LUNs available will be selected as the XenServer Heartbeat SR (see [Heartbeats for Availability](#) for full explanation of this requirement). If there is more than one LUN available, it does not matter which one is selected. A small amount (<5 GB) of space will be allocated from this LUN to provide the HA function.

This reference architecture is designed to support the maintenance of workload within the limits of the unexpected loss of a single host. The 'Failures to Tolerate' setting should be set to 1.

Note:

It is important to be aware of the following characteristics when using HA

- The resource pool must have at least 3 hosts
- Fencing refers to the mechanism that forcibly isolates a host that is considered failed or unsafe, ensuring it cannot access shared resources (especially shared storage) before virtual machines are restarted elsewhere. The larger the resource pool the greater the impact of host fencing is and the more likely it is to occur because of the communication path requirements. Careful consideration should be given in this respect to pools >16 hosts.
- HA should be disabled during host / resource pool maintenance operations to reduce the risk of unexpected outages while these maintenance activities are being undertaken.
- When using HA, care must be taken to ensure that the storage connectivity to the HA statefile SR and the networking is not disrupted. Disruption to these connections can cause hosts to fence unexpectedly to protect the operations of the resource pool. This means that careful consideration must be taken when undertaking maintenance on the infrastructure to ensure connectivity is not lost.

Alternative approaches for Management HA If you are using HA to provide reliable, highly available management connectivity, then the following approaches should also be considered to avoid the limitations listed above.

Implement host monitoring Integrate XenServer hosts with a 3rd-party monitoring solution to detect when a host is unresponsive.

Manually recover management connectivity If the host that is uncontactable is the coordinator and management operations are unavailable, then all the other hosts in the resource pool will enter emergency mode and allow remote connectivity to be established to them to elect one of them as a new pool coordinator. This can be performed using the remote language bindings. Examples below are given for xe CLI and PowerShell.

xe CLI:

```
1 xe -s <server_IP> -u <username> -pw <password> host-is-in-emergency-  
   mode -> Confirm election possibility  
2 xe -s <server_IP> -u <username> -pw <password> pool-emergency-  
   transition-to-master -> Designate current host as new  
   master  
3 xe -s <server_IP> -u <username> -pw <password> pool-recover-slaves  
   -> Reconfigure member servers to new  
   coordinator
```

PowerShell:

```
1 Invoke-XenPool -XenAction EmergencyTransitionToMaster ->  
   Designate current host as new master  
2 Invoke-XenPool -XenAction RecoverSlaves ->  
   Reconfigure member servers to new coordinator
```

VM HA

Workload Balancing (WLB) VM This VM should have the restart priority set to restart to ensure that any WLB functions are not lost after an unexpected outage. It should be in the first start group 0.

Other VMs Depending on the exact usage of XenServer, additional VMs might be configured to use VM HA. The following are suggestions for some use cases.

Decision guidance (VM HA): Use VM HA for workloads where automated restart after an unexpected host outage materially improves service continuity, and where the pool is sized N+1 to absorb the outage. Avoid VM HA for workloads where external brokering/orchestration already provides restart or scale-out behavior (for example, many non-persistent CVAD workloads), or where restart order and dependencies are difficult to control.

Note:

VMs added to the resource pool after HA is configured will default to '**Do Not Restart**'. The HA will need to be reconfigured after VMs are added to the environment if required.

- **PVS and non-persistent MCS VMs:** Do not configure these VMs with any HA. Make sure these are set to **Do Not Restart**. Other VMs should be available to provide immediate services for end users and Citrix Virtual Apps and Desktops will provide power management functions to support the restarting of these VMs either when an end user requires it or through auto scale functions to support the expected demand see [AutoScale](#). This will provide a more optimized outcome for the combined Citrix environment.
- **MCS Persistent VMs:** These should be set to 'Restart'. The configuration being designed should ensure there is enough RAM and CPU available to operate all the VMs in the pool when there is 1 host out of service (for update or maintenance or outage) so this configuration will be possible.
- **General Server Virtualization workloads (any workload not being used to run Citrix VDAs):** These should be set to 'Restart'. The configuration being designed should ensure there is enough RAM and CPU available to operate all the VMs in the pool when there is 1 host out of service (for update or maintenance or outage) so this configuration will be possible. The VMs can be configured to use the start groups and delays as required to ensure the workload restarts as required. This could be used to prioritize workloads that other services depend on such as Active Directory, DHCP and DNS service providers.

Resilience across data centers**Optional:**

Disaster Recovery (DR) across data centers. This reference architecture can support DR if necessary, with additional design and operational considerations.

Decision guidance for DR

- **Use when:** You have a defined business requirement for site resilience, a replication-capable storage platform, and agreed recovery procedures (RPO/RTO) that require VM recovery at a secondary site.
- **Avoid when:** You need frequent workload mobility between sites, or you do not have the operational capacity to test and execute failover/failback processes. For CVAD, do not assume that hypervisor connectivity and certificates will seamlessly fail over without a coordinated CVAD DR design.
- **Operational impact:** DR is not fully automated; failover/failback requires planned runbooks, access to storage replication controls, and regular testing.

The XenServer Disaster Recovery (DR) configuration described here is a mode of resilience that is expected to be an infrequently used operation to support extreme and unplanned outages of infrastructure. This operation is not automated, and recovery requires careful planning and operation. The DR model allows for operation of workloads in a temporary location that may not be an optimal placement but will support critical business operation during a period of outage.

The XenServer DR model is a complex solution and there are other ways to provide protection from site failure. For example, for Citrix workloads where non-persistent VMs are used to provide services, additional VMs can be provisioned into recovery sites to provide temporary services. These options need to be considered before using the XenServer DR function as they may be easier to implement.

It is recommended to avoid using DR for Citrix VDI workloads unless you have a suitable DR solution for CVAD. It is not safe to assume that the same hypervisor connection can fail over to a XenServer DR site, as network and certificate mappings can make this unreliable.

To mitigate the risk of a data center failure and loss of the workload (or critical parts of the workload), a secondary data center location can be used to provide disaster recovery functions. Further documentation is available [here](#).

One or more resource pools in the DR site can be used for other workloads in addition to being the recovery location for workloads, but you must maintain the required spare capacity to enable the additional workloads to be run when required.

There is no requirement for the full workload to be operated in the DR site when a failure event occurs. Partial workloads can be recovered until full operation is restored at the primary site, however there is a requirement that workloads must be recovered from 1 resource pool into a single recovery resource pool. Workloads cannot be distributed to multiple recovery pools.

Important:

The storage replication process that mirrors the data from LUN(s) in the primary site to the DR site must be configured outside of the XenServer tool set and reconfiguration of this must be available to the administrator responsible for failover as they need to break and reset mirroring as part of operations relating to this process (see operations section for details).

When building this deployment:

- All storage used for the VM Virtual disks must be replicated from the primary environment to the backup environment
- All storage used for Pool meta data must be replicated from the primary environment to the backup environment
- The hardware infrastructure at your DR site does not have to match the primary site but must use the same processor vendor. Sufficient memory, CPU, and network capacity must be configured in the standby resource pool to allow all the required failed-over VMs to be recreated and started.


- The XenServer resource pool in the recovery site must be at the same release and patch level or newer than the primary site (see XenServer version management across deployments for details of how to achieve this).

Operations Model

XenServer version management across deployments

XenServer deployments are built from one or more resource pools, across one or more data centers to provide a scalable, robust, and secure deployment capable of delivering performant workloads at scale.

These deployments can provide testing and production environments, and a workflow can be built to roll out changes through these environments.



Optional:

The host testing and workload testing phases are optional. The reference architecture supports operating with or without these phases.

The host testing and workload testing stages do not have to be separate stages; they can be combined. This can be advantageous as it reduces the time to get through changes to production. The XenServer update stream is robust and fully tested, allowing continuous deployment. To maintain security, robustness, and performance, it is important to minimize the time updates take to reach production.

The expectations for each phase are as follows:

- **Early Release Testing:** This phase is intended for early visibility of updates and validation of features. Hosts in this phase are configured to use the early release channel, which typically provides access to new versions 2 weeks before the normal channel. Hosts in this phase must

have outbound connectivity to the internet (either directly or via a proxy server) to access the XenServer Content Distribution Network (CDN).

- **Host Testing:** The purpose of this phase is to test the XenServer software release update processes. Releases are available over the internet from the XenServer CDN or through the offline channel delivery mechanisms.
- **Workload testing/UAT:** This phase enables the XenServer hosts to be tested with a production-like workload to validate scale and resiliency. Releases are available over the internet from the XenServer CDN or through the offline channel delivery mechanisms.
- **Production Deployments:** The hosts running production workloads can be updated following the previous testing phases (directly or through the offline delivery mechanisms).

Note:

Offline updates can be used to ensure that the expected version of XenServer is deployed to the UAT and production environments. It is possible to synchronize updates directly from another resource pool, however, if there are updates synchronized but not yet applied to the source pool, they may be applied to the destination pool, resulting in a higher patch version on the destination pool.

To remain supported, all resource pools need to be updated regularly and not be outdated by more than 6 months.

Driver management with Driver Multi-Versioning (DMV)

XenServer 9 introduces Driver Multi-Versioning (DMV), which delivers approved driver versions through the standard streamed host updates channel, removing the need for separate driver disks in most cases.

DMV allows operators to select from multiple approved driver versions available in the update stream. XenServer presents only driver variants compatible with the hardware on each host. If an incorrect driver is applied, the correct version can be selected and installed without reinstalling the host, though a reboot is required for the change to take effect. XenServer 9 enforces signed-driver validation, blocking unapproved third-party binaries.

For more information, see [Driver multi-versioning](#).

Monitoring the resource pool

When alerts occur in the resource pool for any of the hosts or VMs, these are available in the Alerts section of XenCenter®.

These alerts can also be made available in an external monitoring tool via SNMP integration or through email SMTP integration.

There are 2 types of monitoring to be considered.

- Monitoring for utilization, to ensure that any unexpected capacity issues are identified quickly and can be acted on.
- Monitoring for system events to ensure that infrastructure is operating correctly and that there are no unexpected changes occurring.

These are detailed below to be considered in the design for integration.

Integration to 3rd-party monitoring products can be made through the SNMP feature described [here](#) and can provide alerts that are integrated into those 3rd-party tools.

Note:

When adding a new host to the pool you must reconfigure the SNMP settings to ensure that the new host is included in any SNMP integration (see [constraints](#) for details).

Monitoring for usage There are alerts that can be configured for detecting things like excessive CPU usage, see [Performance Alerts](#) for details.

Below is an example of the data provided in the SNMP traps fired for these events.

```
1 2026-02-19 09:44:35 <UNKNOWN> [UDP:
  [10.71.64.14]:53876->[10.81.143.35]:162]:
2 iso.3.6.1.2.1.1.3.0 = Timeticks: (196826250) 22 days, 18:44:22.50
  iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.4.1.60953.1.10.1
  iso.3.6.1.4.1.60953.1.10.1.1.0 = STRING: "add" iso
  .3.6.1.4.1.60953.1.10.1.2.0 = STRING: "OpaqueRef:eea63153-f96c-25f5
  -17b9-08ce49e7ccf4" iso.3.6.1.4.1.60953.1.10.1.3.0 = STRING: "83
  a9c805-a172-f1b2-55d4-081ec180893b" iso.3.6.1.4.1.60953.1.10.1.4.0 =
  STRING: "ALARM" iso.3.6.1.4.1.60953.1.10.1.5.0 = INTEGER: 3
  iso.3.6.1.4.1.60953.1.10.1.6.0 = STRING: "VM" iso
  .3.6.1.4.1.60953.1.10.1.7.0 = STRING: "681e4064-16cb-4d0c-9c2c
  -887623bc623f" iso.3.6.1.4.1.60953.1.10.1.8.0 = STRING: "20260219T09
  :45:55Z" iso.3.6.1.4.1.60953.1.10.1.9.0 = STRING: "The memory
  required by the control domain on \"Control domain on host: pm-
  demolab-3\" is about 5.7% of its allocated memory. Occasional
  performance degradation can be expected when memory swapping is
  forced to happen.
3 This alarm is set to be triggered when the memory required by the
  control domain is above 5.0% of its allocated memory."
```

Monitoring for system events The system alerts are described [here](#)

Below is an example of the data provided in the SNMP traps fired for these events.

```
1 2026-02-19 09:43:39 <UNKNOWN> [UDP:
  [10.71.64.14]:46982->[10.81.143.35]:162]:
```



```
2 iso.3.6.1.2.1.1.3.0 = Timeticks: (196820650) 22 days, 18:43:26.50
    iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.4.1.60953.1.10.1
iso.3.6.1.4.1.60953.1.10.1.1.0 = STRING: "add" iso
.3.6.1.4.1.60953.1.10.1.2.0 = STRING: "OpaqueRef:ca95e2da-06b7-9aa5-
fa4b-ddf007a50c59" iso.3.6.1.4.1.60953.1.10.1.3.0 = STRING: "525
e459c-51bc-c71c-9c52-dcd747e9a84f" iso.3.6.1.4.1.60953.1.10.1.4.0 =
STRING: "ALL_RUNNING_VMS_IN_ANTI_AFFINITY_GRP_ON_SINGLE_HOST"
iso.3.6.1.4.1.60953.1.10.1.5.0 = INTEGER: 3iso
.3.6.1.4.1.60953.1.10.1.6.0 = STRING: "Pool" iso
.3.6.1.4.1.60953.1.10.1.7.0 = STRING: "419d6735-5d11-56b1-8cdc-2853
ce709a27" iso.3.6.1.4.1.60953.1.10.1.8.0 = STRING: "20260219T09
:44:59Z" iso.3.6.1.4.1.60953.1.10.1.9.0 = STRING: "<body><message>
Breach on VM anti-affinity rules</message><VM_group>3224f9d9-9492-06
f7-3c92-da610cde33a8</VM_group><host>ce2a3c26-6bbc-4d23-8626-
fba28d3dc651</host></body>"
```

Remote Logging The primary purpose of this is to enable support engagement. To ensure the availability of this data during issue resolution, syslog forwarding can be configured. This enables the logs to be retrieved in the event they cannot be retrieved from the host directly, or if there is any reason to suspect logs on the host may have been tampered with.

If this is configured, then this must be configured on all hosts in the pool.

This data can consume significant space. Use and retention of this feature and the data is at your discretion based on internal security policies.

Capacity Planning Model

Sizing the Resource Pool

Each resource pool should contain enough resources so that if 1 host is down the resource pool can operate the full workload. This will enable:

- Planned outage to occur as needed on 1 host at a time
- Resource pool updates to be made without any outage of workload
- Ensure minimal workload outage if 1 host unexpectedly fails

To achieve this the expectation is that all hosts in the resource pool have the same ability to support workload (i.e. same configuration, RAM, CPU). This may result in a higher level of CPU overcommit during outage, but this is expected to be short lived until the outage condition is resolved.

If more than 1 host in a resource pool is lost it will no longer be possible to maintain all workload and there may be a workload outage until hosts are recovered.

Note:

If hosts do not have identical capacity for periods of time during activities like hardware update/refresh, allowance should be made for the failure of the host with the most memory

Resource Pool Capacity calculations In this reference architecture, we have assumed each resource pool is for a specific purpose. As a result, the calculations below assume that all VMs in the resource pool are the same specification.

NUMA optimization XenServer 9 enables NUMA placement optimization by default. When a VM's vCPUs and memory fit within a single NUMA node on the host, workloads achieve lower memory-access latency and more consistent performance. This can also enable higher VM density per host, reducing the need for additional servers.

For best results, size VMs so their memory fits within a single physical NUMA node. If a VM requires more memory than a single node can provide, XenServer spreads the VM across multiple nodes and the optimization benefit is reduced. VM sizing that exceeds NUMA node boundaries should be treated as an exception rather than the norm.

Per-VM NUMA metrics are available to monitor placement and detect affinity violations. See [NUMA optimization in XenServer 9](#) for details.

Resource Pool Capacity calculations Dom0 is sized to 32 GB RAM. However, if PVS accelerator is being used this will need to be increased by at least 5 GB per vdisk to support the acceleration and must be taken into account in these calculations.

Inputs

Term	Definition
host_ram	Total RAM in each Host (in GB)
vm_ram	RAM requirement for each VM (in GB)
host_cpus	Number of vCPUs available in the host (these are logical CPUs so include threads)
vm_cpus	vCPUs required for each VM
hosts_total	Number hosts in pool
vm_disks	Number of disks attached to each VM
vm_snapshots	Expected number of snapshots per VM

Term	Definition
<code>dom0_ram</code>	Memory reserved for Dom0 (32 GB for this reference architecture)
<code>dom0_vcpus</code>	vCPUs reserved for Dom0 (16 for this reference architecture)
<code>wlb_ram</code>	Memory required for WLB appliance (2GB for this reference architecture)
<code>wlb_vcpus</code>	vCPUs required for WLB appliance (2 for this reference architecture)

Outputs

Term	Definition	Formula
<code>vm_host_max</code>	Maximum number VMs running on each host in the pool	$(\text{host\_ram} - \text{dom0\_ram} - \text{wlb\_ram}) / \text{vm\_ram}$
<code>vm_pool_max</code>	Maximum number of VMs that could run in the pool	$\text{vm\_host\_max} * (\text{hosts\_total} - 1)$
<code>srs_total</code>	Number of Storage Repositories needed to host the VMs	$\text{vm\_pool\_max} * (\text{vm\_disks} + \text{vm\_snapshots}) / 1000$
<code>vm_host</code>	Number of VMs actually running on the host	$\text{vm\_pool\_max} / \text{hosts\_total}$ (see note)
<code>host_cpu_overcommit</code>	Number of vCPUs required in the host to run the required workload (these are logical CPUs so include threads)	$((\text{vm\_host} * \text{vm\_cpus}) + \text{dom0\_vcpus} + \text{wlb\_vcpus}) / \text{host\_cpus}$

Note:

During host failure, maintenance, or update operation $\text{vm\_host} = \text{vm\_pool\_max} / (\text{hosts\_total} - 1)$

Worked Example

Term	Value
host_ram	1000 GB (1 TB)
vm_ram	16 GB
host_cpus	2 sockets each with 32 threaded cores = $2 * 32 * 2 = 128$
vm_cpus	4
hosts_total	8
vm_disks	MCS VMs = 1 ID disk, 1 OS disk, 1 MCSIO disk = 3
vm_snapshots	0
dom0_ram	32
dom0_vcpus	16
wlb_ram	2
wlb_vcpus	2

- $vm\_host\_max = (1000 - 32 - 2) / 16 = 60$ max VMs per Host
- $vm\_pool\_max = 60 * (8 - 1) = 420$ max VMs per pool
- $srs\_total = 420 * (3 + 0) / 1000 = 1.26$ SRs (>1 so need 2 SRs spread over 2 LUNS)
- **Normal operation:**
 - $vm\_host = 420 / 8 = 52.5$ VMs per host (round down to 52)
 - $host\_cpu\_overcommit = ((52 * 4) + 16 + 2) / 128 = 1.77$ vCPU per pCPU
- **During single host outage or update/maintenance operation:**
 - $vm\_host = 420 / 7 = 60$ VMs per host
 - $host\_cpu\_overcommit = ((60 * 4) + 16 + 2) / 128 = 2.01$ vCPU per pCPU

Note:

There is no usage of Dynamic Memory Control (DMC) to provide RAM overcommit. XenServer does not support this during normal operation. This is only expected to be used during short periods of planned controlled maintenance when no other option is available to keep required workload running.

If the overcommit is too high, then the number of VMs for the pool will need to be reduced until

it reaches an acceptable point.

System requirements

June 30, 2026

XenServer requires at least two separate physical x86 computers: one to be the XenServer host and the other to run the XenCenter® application or the XenServer Command-Line Interface (CLI). The XenServer host computer is dedicated entirely to the task of running XenServer and hosting VMs, and is not used for other applications.

Warning:

XenServer® supports only drivers and supplemental packs that are provided by us being installed directly in the host's control domain. Drivers provided by third-party websites, including drivers with the same name or version number as those we provide, are not supported.

The following exceptions are supported:

- Software supplied as a supplemental pack and explicitly endorsed by us.
- Drivers that NVIDIA provides to enable vGPU support. For more information, see [NVIDIA vGPU](#).

Other drivers provided by NVIDIA, for example, the Mellanox drivers, are not supported with XenServer unless distributed by us.

To run XenCenter use any general-purpose Windows system that satisfies the hardware requirements. This Windows system can be used to run other applications.

When you install XenCenter on this system, the XenServer CLI is also installed. A standalone remote XenServer CLI can be installed on any RPM-based Linux distribution. For more information, see [Command-line interface](#).

XenServer host system requirements

Although XenServer is usually deployed on server-class hardware, XenServer is also compatible with many models of workstations and laptops. For more information, see the [Hardware Compatibility List \(HCL\)](#).

The following section describes the recommended XenServer hardware specifications.

The XenServer host must be a 64-bit x86 server-class machine devoted to hosting VMs. XenServer creates an optimized and hardened Linux partition with a Xen-enabled kernel. This kernel controls the interaction between the virtualized devices seen by VMs and the physical hardware.

XenServer can use:

- Up to 6 TB of RAM
- Up to 16 physical NICs
- Up to 960 logical processors per host.

Note:

The maximum number of logical processors supported differs by CPU. For more information, see the [Hardware Compatibility List \(HCL\)](#).

The system requirements for the XenServer host are:

CPUs

One or more 64-bit x86 CPUs, 1.5 GHz minimum, 2 GHz or faster multicore CPU recommended.

Note:

Ensure that you enable hardware support for virtualization on the XenServer host. Virtualization support is an option in your system firmware. It is possible that your hardware might have virtualization support disabled. For more information, see your server documentation.

Supported CPU Architecture Requires CPUs which support the x86-64-v3 microarchitecture.

Supported GPU Architecture

Note:

vGPU support for Nvidia GPU cards will only be available at the full GA release version.

RAM

2 GB minimum, 4 GB or more recommended

Disk space

- Locally attached storage with 46 GB of disk space minimum, 70 GB of disk space recommended
- SAN via HBA (not through software) when installing with multipath boot from SAN.

For a detailed list of compatible storage solutions, see the [Hardware Compatibility List \(HCL\)](#).

Network

100 Mbit/s or faster NIC. One or more Gb, or 10 Gb NICs is recommended for faster export/import data transfers and VM live migration.

We recommend that you use multiple NICs for redundancy. The configuration of NICs differs depending on the storage type. For more information, see the vendor documentation.

XenServer requires an IPv4 network for management and storage traffic.

Notes:

- Ensure that the time setting on your server is set to the current time in UTC.
- In some support cases, serial console access is required for debug purposes. When setting up the XenServer configuration, we recommend that you configure serial console access. For hosts that do not have physical serial port or where suitable physical infrastructure is not available, investigate whether you can configure an embedded management device. For more information about setting up serial console access, see [CTX228930 - How to Configure Serial Console Access on XenServer](#).

XenCenter system requirements

XenCenter has the following system requirements:

- **Operating System:**

- Windows 10
- Windows 11
- Windows Server 2016
- Windows Server 2019
- Windows Server 2022
- Windows Server 2025

- **.NET Framework:** Version 4.8
- **CPU Speed:** 750 MHz minimum, 1 GHz or faster recommended
- **RAM:** 1 GB minimum, 2 GB or more recommended
- **Disk Space:** 100 MB minimum
- **Network:** 100 Mbit/s or faster NIC
- **Screen Resolution:** 1024x768 pixels, minimum

If you want XenCenter to be able to launch an external SSH console that connects to your XenServer host, install one of the following applications on the system:

- PuTTY
- OpenSSH (installed by default on some Windows operating systems)

For more information, see [Configure XenCenter to use an external SSH console](#).

Supported guest operating systems

For a list of supported VM operating systems, see [Guest operating system support](#).

Pool requirements

A resource pool is a homogeneous or heterogeneous aggregate of one or more hosts, up to a maximum of 32. Before you create a pool or join a host to an existing pool, ensure that all hosts in the pool meet the following requirements.

Hardware requirements

All of the servers in a XenServer resource pool must have broadly compatible CPUs, that is:

- The CPU vendor (Intel, AMD) must be the same on all CPUs on all servers.
- All CPUs must have virtualization enabled.

Other requirements

In addition to the hardware prerequisites identified previously, there are some other configuration prerequisites for a host joining a pool:

- It must have a consistent IP address (a static IP address on the host or a static DHCP lease). This requirement also applies to the servers providing shared NFS or iSCSI storage.
- Its system clock must be synchronized to the pool coordinator (for example, through NTP).
- It cannot be a member of an existing resource pool.
- It cannot have any running or suspended VMs or any active operations in progress on its VMs, such as shutting down or exporting. Shut down all VMs on the host before adding it to a pool.
- It cannot have any shared storage already configured.

- It cannot have a bonded management interface. Reconfigure the management interface and move it on to a physical NIC before adding the host to the pool. After the host has joined the pool, you can reconfigure the management interface again.
- It must be running the same version of XenServer, at the same patch level, as hosts already in the pool.
- It must be configured with the same supplemental packs as the hosts already in the pool. Supplemental packs are used to install add-on software into the XenServer control domain, dom0. To prevent an inconsistent user experience across a pool, all hosts in the pool must have the same supplemental packs at the same revision installed.
- It must have the same XenServer license as the hosts already in the pool. You can change the license of any pool members after joining the pool. The host with the lowest license determines the features available to all members in the pool.

XenServer hosts in resource pools can contain different numbers of physical network interfaces and have local storage repositories of varying size.

Note:

Servers providing shared NFS or iSCSI storage for the pool must have a static IP address or be DNS addressable.

Homogeneous pools

A homogeneous resource pool is an aggregate of servers with identical CPUs. CPUs on a server joining a homogeneous resource pool must have the same vendor, model, and features as the CPUs on servers already in the pool.

Heterogeneous pools

Heterogeneous pool creation is made possible by using technologies in Intel (FlexMigration) and AMD (Extended Migration) CPUs that provide CPU *masking* or *leveling*. These features allow a CPU to be configured to *appear* as providing a different make, model, or feature set than it actually does. These capabilities enable you to create pools of hosts with different CPUs but still safely support live migrations.

For information about creating heterogeneous pools, see [Hosts and resource pools](#).

Configuration limits

June 30, 2026

Use the following configuration limits as a guideline when selecting and configuring your virtual and physical environment for XenServer. The following tested and recommended configuration limits are fully supported for XenServer.

- Virtual machine limits
- XenServer host limits
- Resource pool limits

Factors such as hardware and environment can affect the limitations listed below. More information about supported hardware can be found on the [Hardware Compatibility List](#). Consult your hardware manufacturers' documented limits to ensure that you do not exceed the supported configuration limits for your environment.

Virtual machine (VM) limits

Item	Limit
Compute	
Virtual CPUs per VM (Linux)	32/64 (see note 1)
Virtual CPUs per VM (Windows)	32/64 (see note 1)
Memory	
RAM per VM	1.5 TiB (see note 2)
Storage	
Virtual Disk Images (VDI) (including CD-ROM) per VM	241 (see note 3)
Virtual CD-ROM drives per VM	1
Virtual Disk Size (NFS)	2040 GiB
Virtual Disk Size (LVM)	2040 GiB
Virtual Disk Size (XFS)	16 TiB

Item	Limit
Virtual Disk Size (GFS2)	16 TiB

Networking

Virtual NICs per VM	7 (see note 4)

Graphics Capability

vGPUs per VM	8
Passed through GPUs per VM	1

Devices

Pass-through USB devices	1 (see note 5)

Notes:

1. Consult your guest OS documentation to ensure that you do not exceed the supported limits. As of yet, Red Hat Enterprise Linux 8 and derivatives do not support more than 32 vCPUs. Though the limit is 64, we recommend setting the limit to 32 if your VMs might not be trustworthy or if you want to prevent a potential impact on system availability.
2. The maximum amount of physical memory addressable by your operating system varies. Setting the memory to a level greater than the operating system supported limit may lead to performance issues within your guest.
3. The maximum number of VDIs supported depends on the guest operating system. Consult your guest operating system documentation to ensure that you do not exceed the supported limits.
4. Several guest operating systems have a lower limit, other guests require installation of the XenServer VM Tools to achieve this limit.
5. If you assign multiple USB passthrough devices to a VM, the VM might fail to start. See [Known Issues](#).

XenServer host limits

Item	Limit
Compute	
Logical processors per host	1024 (see note 1)
Concurrent VMs per host	200 (see note 2)
Concurrent protected VMs per host with HA enabled	500
Virtual GPU VMs per host	128 (see note 3)
Memory	
RAM per host	6 TB
Storage	
Concurrent active virtual disks per host	2048 (see note 4)
Storage repositories per host (NFS)	400
Attached VDIs per SR (all types)	600
Networking	
Physical NICs per host	16
Physical NICs per network bond	4
Virtual NICs per host	512
VLANs per host	800
Network Bonds per host	4
Graphics Capability	
GPUs per host	8 (see note 5)

Notes:

1. The maximum number of logical physical processors supported differs by CPU. For more information, see the [Hardware Compatibility List](#).
2. The maximum number of VMs per host supported depends on VM workload, system load,

network configuration, and certain environmental factors. We reserve the right to determine what specific environmental factors affect the maximum limit at which a system can function. When using the PVS Accelerator, we recommend allocating at least 16 GB RAM to the Control Domain. For information about configuring Dom0 memory, see [Change the amount of memory allocated to the control domain](#).

3. For NVIDIA vGPU, 128 vGPU accelerated VMs per host with 4xM60 cards (4x32=128 VMs), or 2xM10 cards (2x64=128 VMs).
4. The number of concurrent active virtual disks per host is also constrained by the number of SRs you have attached to the host and the number of attached VDIs that are allowed for each SR (600). For more information, see the “Attached VDIs per SR” entry in the Resource pool limits.
5. This figure might change. For the current supported limits, see the [Hardware Compatibility List](#).

Resource pool limits

Item	Limit
Compute	
VMs per resource pool	2400
Hosts per resource pool	32 (see note 1)
Networking	
VLANs per resource pool	800
Disaster recovery	
Integrated site recovery storage repositories per resource pool	8
Storage	
Paths to a LUN	16
Multipathed LUNs per host	150 (see note 2)
Multipathed LUNs per host (used by storage repositories)	150 (see note 2)

Item	Limit
VDIs per SR (NFS, SMB, EXT, XFS, GFS2)	20000 (see note 3)
VDIs per SR (LVM)	1000 (see note 3)
Storage repositories per pool (NFS)	400
Storage repositories per pool (GFS2)	62
Maximum file system size (GFS2)	100 TiB

Storage live migration

(non-CDROM) VDIs per VM	6
Snapshots per VM	1
Concurrent transfers	3

XenCenter

Concurrent operations per pool	25

Notes:

1. Clustered pools that use GFS2 storage support a maximum of 16 hosts in the resource pool. For all other cases, 32 is the maximum supported pool size. However, it is not the optimal pool size for most environments and workloads. For more information, see [Recommended pool size](#)
2. When HA is enabled, we recommend increasing the default timeout to at least 120 seconds when more than 30 multipathed LUNs are present on a host. For information about increasing the HA timeout, see [Configure high availability timeout](#).
3. The VDI count limit includes VDIs used for internal purposes such as snapshot management. Each snapshot is internally represented as 2 VDIs: one is the snapshot and one is the shared parent with the active writeable VDI.

Workload Balancing limits

Item	Limit
Pools managed per Workload Balancing virtual appliance	100 (see note 1)

Item	Limit
VMs managed per Workload Balancing virtual appliance	1000 (see note 2)

Notes:

1. Ensure that the virtual appliance has sufficient resources (vCPU, memory, disk size) to manage large numbers of pools and VMs.
2. If a single pool has a large number of VMs (for example, more than 400 VMs), we recommend that you use one Workload Balancing virtual appliance just for that pool.

Hardware drivers

February 19, 2026

We collaborate with partner organizations to provide drivers and support for a wide range of hardware. For more information, see [the Hardware Compatibility List](#).

To support this hardware, your installation of XenServer 9 includes third-party drivers that have been certified as compatible with XenServer.

Updates to drivers

We regularly deliver updated versions of these drivers, which can enable new hardware or resolve issues with existing hardware. Most driver updates are delivered through the updates mechanism. For more information, see [Update your XenServer hosts](#).

Even though we distribute the drivers and their source code for our customers to use, the hardware vendor owns the driver source files.

When updating drivers either from the updates mechanism, review the information in the [Hardware Compatibility List \(HCL\)](#) and your hardware vendor documentation to ensure that your hardware is installed with a firmware version that is compatible with the updated driver.

Support for drivers

XenServer supports only drivers that are delivered in-box with the product or are downloaded from <https://support.citrix.com>. Drivers provided by third-party websites, including drivers with the same name or version number as those drivers provided by us, are not supported.

Note:

The only exception to this restriction are the drivers that NVIDIA provides to enable vGPU support. For more information, see [NVIDIA vGPU](#).

Other drivers provided by NVIDIA, for example, the Mellanox drivers, are only supported with XenServer® when distributed by us.

Do not download drivers from your hardware vendor website, even if the driver has the same version number as the one provided by XenServer. These drivers are not supported.

Before a driver can be supported with XenServer, it must be certified with us and released through one of the approved mechanisms. This certification process ensures that the driver is of a format required to be installable in a XenServer environment and that it is compatible with XenServer 9.

What if a driver I need isn't supported?

If your hardware vendor recommends that you install a specific driver version that is not available in-box or on the <https://support.citrix.com> website, request that the vendor contacts us to certify this version of the driver with XenServer.

We provide the vendors with certification kits that they can use to test updated versions of their drivers that are required by the shared customer base of XenServer and the hardware vendor. After the vendor provides us with the certification test results, we validate that those results show no issues or regressions in the updated version of the driver. The driver version is now certified with XenServer and we publish the driver through our regular updates.

For more information about the certification process the vendor must follow, see the article [Hardware Compatibility List explained](#).

Guest operating system support

July 6, 2026

When installing VMs and allocating resources such as memory and disk space, follow the guidelines of the operating system and any relevant applications.

Operating System	Minimum RAM	Maximum RAM	Minimum Disk Space
Windows 10 (64-bit)	2 GB	1.5 TB	32 GB (40 GB or more recommended)

Operating System	Minimum RAM	Maximum RAM	Minimum Disk Space
Windows 11	4 GB	1.5 TB	32 GB (40 GB or more recommended)
Windows Server 2016, Windows Server Core 2016	1 GB	1.5 TB	32 GB (40 GB or more recommended)
Windows Server 2019, Windows Server Core 2019	1 GB	1.5 TB	32 GB (40 GB or more recommended)
Windows Server 2022, Windows Server Core 2022	1 GB	1.5 TB	32 GB (40 GB or more recommended)
Windows Server 2025, Windows Server Core 2025 (see note 2)	2 GB	1.5 TB	64 GB (64 GB or more recommended)
CentOS Stream 9 (64-bit) (see note 1)	2 GB	1.5 TB	10 GB
CentOS Stream 10 (64-bit) (see note 1)	2 GB	1.5 TB	10 GB
Red Hat Enterprise Linux 8 (64-bit)	2 GB	1.5 TB	10 GB
Red Hat Enterprise Linux 9 (64-bit) (see note 1)	2 GB	1.5 TB	10 GB
Red Hat Enterprise Linux 10 (64-bit) (preview) (see note 1)	2 GB	1.5 TB	10 GB
SUSE Linux Enterprise Server 12 SP5 (64-bit)	1 GB	1.5 TB	8 GB
SUSE Linux Enterprise Server 15 SP5, 15 SP6 (64-bit)	1 GB	1.5 TB	8 GB
SUSE Linux Enterprise Desktop 15 SP5, 15 SP6 (64-bit)	1 GB	1.5 TB	8 GB
Oracle Linux 8 (64-bit)	2 GB	1.5 TB	10 GB
Debian Bullseye 11 (64-bit)	512 MB	1.5 TB	10 GB

Operating System	Minimum RAM	Maximum RAM	Minimum Disk Space
Debian Bookworm 12 (64-bit) (see note 1)	1 GB	1.5 TB	10 GB
Ubuntu 22.04 (64-bit)	1 GB	1.5 TB	10 GB
Ubuntu 24.04 (64-bit)	1 GB	1.5 TB	10 GB
Goroom 2 (64-bit)	1 GB	1.5 TB	10 GB
Rocky Linux 8 (64-bit)	1 GB	1.5 TB	10 GB
Rocky Linux 9 (64-bit) (see note 1)	2 GB	1.5 TB	15 GB

Notes:

1. Customers who wish to use this guest OS must also install XenServer VM Tools for Linux v8.3.1-1 or later, available to download from the [XenServer product downloads page](#).
 2. When installing Windows Server 2025 VMs, having two disks attached might cause the VM to hang during boot. To avoid this issue, complete the installation with a single disk. Next, install VM Tools for Windows on the VM and reboot to finalize the installation. Once complete, you can safely attach the second disk.
- XenServer® VM Tools for Linux is only supported on the Linux guest operating systems listed above.
 - All supported operating systems run in HVM mode.
 - Individual versions of the operating systems can also impose their own maximum limits on the amount of memory supported (for example, for licensing reasons).
 - When configuring guest memory, do not to exceed the maximum amount of physical memory that your operating system can address. Setting a memory maximum that is greater than the operating system supported limit might lead to stability problems within your guest.
 - To create a VM of a newer minor version of RHEL than is listed in the preceding table, use the following method:
 - Install the VM from the latest supported media for the major version
 - Use `yum update` to update the VM to the newer minor version

This approach also applies to RHEL-based operating systems such as CentOS and Oracle Linux.

- XenServer supports all SKUs (editions) for the listed versions of Windows.

Long-term guest support

XenServer includes a long-term guest support (LTS) policy for Linux VMs. The LTS policy enables you to consume minor version updates by one of the following methods:

- Installing from new guest media
- Upgrading from an existing supported guest

Out-of-support operating systems

The list of supported guest operating systems can contain operating systems that were supported by their vendors at the time this version of XenServer was released, but are now no longer supported by their vendors.

We no longer offer support for these operating systems (even if they remain listed in the table of supported guests or their templates remain available on your XenServer hosts). While attempting to address and resolve a reported issue, we assess if the issue directly relates to an out-of-support operating system on a VM. To assist in making that determination, we might ask you to attempt to reproduce an issue using a supported version of the guest operating system. If the issue seems to be related to the out-of-support operating system, we will not investigate the issue further.

Note:

Windows versions that are supported by Microsoft as part of an LTSB branch are supported by XenServer.

Windows versions that are out of support, but part of an Extended Security Updates (ESU) agreement are not supported by XenServer.

Feature support

Some XenServer features are supported only by a subset of the guest operating systems. The following table lists whether an operating supports a particular feature.

Operating System	Conversion		Dynamic Mem-ory		vGPU live mi-gration	vGPU	GPU passthrough	PCI passthrough	USB passthrough	SR-IOV passthrough
	Man-ager	UEFI boot	Control							
Windows 10	YES	YES	YES	YES	YES	YES	YES	YES	YES	
Windows 11	YES	YES	YES	YES	YES	YES	YES	YES	YES	

Operating System	Conversion		Dynamic Mem-ory Control		vGPU live mi-gration	GPU passthrough	PCI passthrough	USB passthrough	SR-IOV passthrough
	Man-ager	UEFI boot		vGPU					
Windows Server 2016, Win-dows Server Core 2016	YES	YES	YES	YES	YES	YES	YES	YES	YES
Windows Server 2019, Win-dows Server Core 2019	YES	YES	YES	YES	YES	YES	YES	YES	YES
Windows Server 2022, Win-dows Server Core 2022	YES	YES	YES	YES	YES	YES	YES	YES	YES
Windows Server 2025, Win-dows Server Core 2025	(Pre-view)								

Operating System	Conversion Manager	UEFI boot	Dynamic Memory Control		vGPU live migration	GPU passthrough	PCI passthrough	USB passthrough	SR-IOV passthrough
				vGPU					
CentOS Stream 9 (64-bit) (see note 1)				YES		YES	YES	YES	
CentOS Stream 10 (64-bit)		YES (see note 2)		YES		YES	YES	YES	YES
Red Hat Enterprise Linux 8 (64-bit)	YES	YES		YES		YES	YES	YES	YES
Red Hat Enterprise Linux 9 (64-bit) (see note 1)	YES	YES		YES		YES	YES	YES	YES
Red Hat Enterprise Linux 10 (64-bit) (pre-view) (see note 1)		YES		YES		YES	YES	YES	YES

Operating System	Conversion Manager	UEFI boot	Dynamic Mem-ory Control		vGPU live mi-gration	GPU passthrough	PCI passthrough	USB passthrough	SR-IOV passthrough
				vGPU					
SUSE Linux Enterprise Server 12 SP5 (64-bit)			YES				YES	YES	
SUSE Linux Enterprise Server 15 SP5, 15 SP6 (64-bit)		YES	YES				YES	YES	
SUSE Linux Enterprise Desktop 15 SP5, 15 SP6 (64-bit)		YES	YES				YES	YES	YES
Oracle Linux 8 (64-bit)		YES	YES				YES	YES	YES
Debian Bulls-eye 11 (64-bit)			YES						

Operating System	Conversion		Dynamic Mem-ory Control	vGPU	vGPU live mi-gration	GPU passthrough	PCI passthrough	USB passthrough	SR-IOV passthrough
	Man-ager	UEFI boot							
Debian Bookworm 12 (64-bit) (see note 1)		YES	YES				YES	YES	YES
Ubuntu 22.04 (64-bit)	YES	YES	YES	YES		YES	YES	YES	
Ubuntu 24.04 (64-bit)	YES	YES	YES	YES		YES		YES	YES
Goroom 2 (64-bit)			YES				YES	YES	YES
Rocky Linux 8 (64-bit)	YES	YES		YES		YES	YES	YES	YES
Rocky Linux 9 (64-bit) (see note 1)	YES	YES		YES		YES	YES	YES	YES

Notes:

1. All UEFI boot or UEFI Secure Boot VMs allow you to attach a vTPM.
2. CentOS 10 Stream supports UEFI boot, but Secure Boot is not currently supported due to an upstream issue.

Connectivity requirements

April 28, 2026

This article provides an overview of domains and common ports that are used by XenServer® components and must be considered as part of the networking architecture, especially if communication traffic traverses network components such as firewalls or proxy servers where ports must be opened or domains added to an allow list to ensure communication flow.

External domains accessed by XenServer product components

Depending on your deployment and requirements, configure your firewall to enable these XenServer components to access the listed domains.

XenServer hosts

Your XenServer hosts access the following domains:

Domain	Port	Direction	Details
repo.ops.xenserver.com	443	Outbound	The XenServer pool coordinator downloads available updates for XenServer 9 from this location. For more information, see Updates .
repo-src.ops.xenserver.com	443	Outbound	The XenServer pool coordinator downloads the source files for XenServer 9 updates from this location. For more information, see Updates .

Domain	Port	Direction	Details
telemetry.ops.xenserver.com	443	Outbound	The XenServer pool coordinator gathers telemetry data and uploads it regularly to this location. For more information, see Telemetry .

When configuring your XenServer pools to receive updates, you can configure a proxy server for the pool coordinator to use to download the updates. For more information, see [Configure updates for your pool](#).

XenCenter®

The XenCenter management console accesses the following domains:

Domain	Port	Direction	Details
updates.ops.xenserver.com	443	Outbound	XenCenter polls information on this site to see whether updates are available for XenCenter and for XenServer 9 hosts. For more information, see Update your XenServer hosts

You can configure a proxy server that XenCenter goes through to check for and download updates. For more information, see [Proxy server](#).

Windows VMs

If you have set up your Windows VMs to receive updates to the XenServer VM Tools Management Agent, your Windows VM accesses the following domains:

Domain	Port	Direction	Details
<code>pvupdates.vmd.citrix.com</code>	443	Outbound	The XenServer VM Tools for Windows poll information on this site to see whether updates are available for the Management Agent.
<code>downloadns.citrix.com.edgesuite.net</code>	443	Outbound	The XenServer VM Tools for Windows download the installer files for the Management Agent from this location.

If you don't want your Windows VM to access these domains, you can redirect Management Agent updates to an internal web server. For more information, see [Redirect the Management Agent updates](#).

Communication ports used by XenServer product components

The ports listed in the following table are the common ports that are used by XenServer components. Not all ports need to be open, depending on your deployment and requirements.

Source	Destination	Type	Port	Details
XenServer hosts	XenServer hosts	TCP	80, 443	Intra-host communication between members of a resource pool using the management API
		UDP	694	High availability (non-clustering) network heartbeat

Source	Destination	Type	Port	Details
	Citrix License Server	TCP	27000	Handles initial connection for license requests
		TCP	7279	Check-in/check-out of licenses
		TCP	8083	Upload telemetry data and connect to cloud-based licensing
	NTP Service	TCP, UDP	123	Time Synchronization
	DNS Service	TCP, UDP	53	DNS Lookups
	Domain Controller	TCP, UDP	389	LDAP (for Active Directory user authentication)
		TCP	636	LDAP over SSL (LDAPS)
	FileServer (with SMB storage)	TCP, UDP	139	ISOStore:NetBIOSSessionService
		TCP, UDP	445	ISOStore:Microsoft-DS
	SAN Controller	TCP	3260	iSCSI Storage
	NAS Head/File Server	TCP	2049	NFSv4 Storage
		TCP, UDP	2049	NFSv3 Storage. TCP is the default
		TCP, UDP	111	NFSv3 Storage - connection to rpcbind
		TCP, UDP	Dynamic	NFSv3 Storage - a dynamic set of ports chosen by the filer
	Syslog	UDP	514	Sends data to a central location for collation

Source	Destination	Type	Port	Details
	GFS2 Clustering	TCP	8892, 8896, 21064	Communication between all pool members in a GFS2 clustered pool.
		UDP	5404, 5405	
	Workload Balancing virtual appliance	TCP	8012	By default, the Workload Balancing server uses 8012. However, if you specify a different port during Workload Balancing set up, ensure that communication is allowed on that port.
XenCenter	XenServer hosts	TCP	22	SSH
		TCP	443	Management using the management API
	Virtual Machine	TCP	5900	VNC for Linux VMs
		TCP	3389	RDP for Windows VMs
Workload Balancing virtual appliance	XenServer hosts	TCP	443	XenServer hosts use port 443 for Workload Balancing to gather metric data.

Source	Destination	Type	Port	Details
Other clients	XenServer hosts	TCP	80, 443	Any client that uses the management API to communicate with XenServer hosts
		UDP	161	SNMP monitoring of XenServer hosts (when configured)
		TCP	5666	NRPE monitoring of XenServer hosts (when configured)

XenServer interoperates with various Citrix products. For more information about the ports these products use, see [Communication ports used by Citrix](#).

Note:

- To improve security, you can close TCP port 80 on the management interface of XenServer hosts. For more information about how to close port 80, see [Restrict use of port 80](#).
- If FQDN is used instead of IP as a resource, then make sure it is resolvable.

Active Directory integration

If you use Active Directory (AD) in your environment, ensure that the following ports are open for connections outbound from XenServer to the AD domain controllers.

Port	Protocol	Use
53	UDP/TCP	DNS
88	UDP/TCP	Kerberos 5
123	UDP	NTP
135	TCP	RPC Endpoint Mapper
137	UDP	NetBIOS Name Service

Port	Protocol	Use
139	TCP	NetBIOS Session (SMB)
389	UDP/TCP	LDAP
445	TCP	SMB over TCP
464	UDP/TCP	Machine password changes
636	UDP/TCP	LDAP over SSL
3268	TCP	Global Catalog Search
49152-65535	TCP	RPC dynamic connections

For more information, see [Active Directory integration](#)

Citrix Provisioning™ Services

If you use Citrix Provisioning Services in your environment, ensure that the following firewall ports can be accessed:

Port	Protocol	Use
6901, 6902, 6905	UDP	Provisioning server outbound communication (packets destined for the target device)
6910	UDP	Target device logon with Citrix Provisioning Services
6901	UDP	Configurable target device port. The default port is 6901.
6910–6930	UDP	Configurable server port range. The default range is 6910–6930.

For more information, see [Citrix Provisioning Services](#) and [Communication ports used by Citrix](#).

Technical overview

June 30, 2026

XenServer is an industry leading platform for cost-effective desktop, server, and cloud virtualization infrastructures. XenServer enables organizations of any size or type to consolidate and transform compute resources into virtual workloads for today's data center requirements. Meanwhile, it ensures a seamless pathway for moving workloads to the cloud.

The key features of XenServer® are:

- Consolidating multiple virtual machines (VMs) onto a physical server
- Reducing the number of separate disk images to be managed
- Allowing for easy integration with existing networking and storage infrastructures
- Enabling you to schedule zero downtime maintenance by live migrating VMs between XenServer hosts
- Assuring availability of VMs by using high availability to configure policies that restart VMs on another host in case one fails
- Increasing portability of VM images, as one VM image works on a range of deployment infrastructures


Virtualization and hypervisor

Virtualization, or to be more specific, hardware virtualization, is a method of running multiple independent VMs on a single physical computer. Software run on these virtual machines is separated from the underlying hardware resources. It's a way of fully utilizing the physical resources available in modern powerful servers, which reduces the total cost of ownership (TCO) for server deployments.

A hypervisor is the basic abstraction layer of software. The hypervisor performs low-level tasks such as CPU scheduling and is responsible for memory isolation for resident VMs. The hypervisor abstracts the hardware for the VMs. The hypervisor has no knowledge of networking, external storage devices, video, and so on.

Key components

This section gives you a high-level understanding of how XenServer works. See the following illustration for the key components of XenServer:



Architecture overview

Hardware

The hardware layer contains the physical server components, such as CPU, memory, network, and disk drives.

You need an Intel VT or AMD-V 64-bit x86-based system with one or more CPUs to run all supported guest operating systems. For more information about XenServer host system requirements, see [System requirements](#).

For a complete list of XenServer certified hardware and systems, see the [Hardware Compatibility List \(HCL\)](#).

Xen® Hypervisor

The Xen Project hypervisor is an open-source type-1 or bare-metal hypervisor. It allows many instances of an operating system or different operating systems to run in parallel on a single machine (or host). Xen hypervisor is used as the basis for many different commercial and open-source applications, such as: server virtualization, Infrastructure as a Service (IaaS), desktop virtualization, security applications, embedded, and hardware appliances.

XenServer is based on the Xen Project hypervisor and on top of that we provide extra features and support. XenServer uses version 4.13.4 of the Xen hypervisor.

Control domain

The **Control Domain**, also called Domain 0, or dom0, is a secure, privileged Linux VM that runs the XenServer management toolstack known as XAPI. In XenServer 9, this Linux VM is running using a custom XenServer OS distribution. Besides providing XenServer management functions, dom0 also runs the physical device drivers for networking, storage, and so on. The control domain can talk to the hypervisor to instruct it to start or stop guest VMs.

Note that whilst the Control Domain is a Linux VM, installing additional code or packages or modifying the configuration is not supported unless explicitly instructed by official XenServer documentation or under the guidance of the Support Team.

Toolstack The **Toolstack**, or XAPI is the software stack that controls VM lifecycle operations, host and VM networking, VM storage, and user authentication. It also allows the management of XenServer resource pools.

XAPI provides the publicly documented management API, which is used by all tools that manage VMs, and resource pools. For more information, see the [XenServer Management API](#).

Guest domain (VMs)

Guest domains are user-created virtual machines that request resources from dom0. For a detailed list of the supported distributions, see [Supported Guests, Virtual Memory, and Disk Size Limits](#).

Full virtualization Full virtualization, or hardware-assisted virtualization uses virtualization extensions from the host CPU to virtualize guests. Fully virtualized guests do not require any kernel support. The guest is called a hardware virtual machine (HVM). HVM requires Intel VT or AMD-V hardware extensions for memory and privileged operations. XenServer uses Quick Emulator (QEMU) to emulate PC hardware, including BIOS, IDE disk controller, VGA graphic adaptor, USB controller, network adapter, and so on. To improve the performance of hardware-sensitive operations like disk or network access, HVM guests are installed with the XenServer tools. For more information, see [PV on HVM](#).

HVM is commonly used when virtualizing an operating system such as Microsoft Windows where it is impossible to modify the kernel to make it virtualization aware.

PV on HVM PV on HVM is a mixture of paravirtualization and full hardware virtualization. The primary goal is to boost performance of HVM guests by using specially optimized paravirtualized drivers.

This mode allows you to take advantage of the x86 virtual container technologies in newer processors for improved performance. Network and storage access from these guests still operate in PV mode, using drivers built in to the kernels.

Windows and Linux distributions are available in PV on HVM mode in XenServer. For a list of supported distributions using PV on HVM, see [Guest Operating System Support](#).

XenServer VM Tools XenServer VM Tools (formerly Citrix VM Tools or XenServer PV Tools) provide high performance I/O services without the overhead of traditional device emulation.

- XenServer VM Tools for Windows consist of I/O drivers (also known as paravirtualized drivers or PV drivers) and the Management Agent.

The I/O drivers contain front-end storage and network drivers, and low-level management interfaces. These drivers replace the emulated devices and provide high-speed transport between VMs and XenServer product family software.

The Management Agent, also known as the guest agent, is responsible for high-level virtual machine management features. It provides full functionality to XenCenter® (for Windows VMs).

XenServer VM Tools for Windows must be installed on each Windows VM for the VM to have a fully supported configuration. A VM functions without the XenServer VM Tools for Windows, but performance will be significantly hampered when the I/O drivers (PV drivers) are not installed.

- XenServer VM Tools for Linux contain a guest agent that provides extra information about the VM to the host. Install the guest agent on each Linux VM to enable Dynamic Memory Control (DMC).

Note:

You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Red Hat Enterprise Linux 10 (preview), Rocky Linux 8, Rocky Linux 9, CentOS Stream 9, or CentOS Stream 10 VMs as these operating systems do not support memory ballooning with the Xen hypervisor.


For more information, see [XenServer VM Tools](#).

Key concepts

Resource pool

XenServer allows you to manage multiple hosts and their connected shared storage as a single entity by using resource pools. Resource pools enable you to move and run virtual machines on different XenServer hosts. They also allow all hosts to share a common framework for network and storage. A

pool can contain up to 32 hosts running the same version of XenServer software, at the same patch level, and with broadly compatible hardware. For more information, see [Hosts and resource pools](#).




Resource pool overview

XenServer resource pool adopts a primary/secondaries architecture, implemented by XAPI. XAPI calls are forwarded from the pool coordinator (the primary) to pool members (the secondaries). Pool members make DB RPCs against the pool coordinator. The pool coordinator is responsible for coordination and locking resources within the pool, and processes all control operations. Pool members talk to the pool coordinator through HTTP and XMLRPC, but they can talk to each other (over the same channel) through mirror disks (storage migration)

Storage repository

XenServer storage targets are called storage repositories (SRs). A storage repository stores Virtual Disk Images (VDIs), which contains the contents of a virtual disk.

SRs are flexible, with built-in support for SATA, SCSI, NVMe, and SAS drives that are locally connected, and iSCSI, NFS, SAS, SMB, and Fibre Channel remotely connected. The SR and VDI abstractions allow advanced storage features such as thin provisioning, VDI snapshots, and fast cloning to be exposed on storage targets that support them.



Storage overview


Each XenServer host can use multiple SRs and different SR types simultaneously. These SRs can be shared between hosts or dedicated to particular hosts. Shared storage is pooled between multiple hosts within a defined resource pool. A shared SR must be network-accessible to each host in the pool. All hosts in a single resource pool must have at least one shared SR. Shared storage cannot be shared between multiple pools.

For more information about how to operate with SRs, see [Configure storage](#).

Networking

On an architecture level, there are three types of server-side software objects to represent networking entities. These objects are:

- A **PIF**, which is a software object used within dom0 and represents a physical NIC on a server. PIF objects have a name and description, a UUID, the parameters of the NIC that they represent, and the network and host they are connected to.
- A **VIF**, which is a software object used within dom0 and represents a virtual NIC on a virtual machine. VIF objects have a name and description, a UUID, and the network and VM they are connected to.
- A **network**, which is a virtual Ethernet switch on a host used to route network traffic on a network host. Network objects have a name and description, a UUID, and the collection of VIFs and PIFs connected to them.



Networking overview

XenServer management APIs allow following operations:

- Configuration of networking options
- Control over the NIC to be used for management operations
- Creation of advanced networking features such as VLANs and NIC bonds

For more information about how to manage networks on XenServer, see [Networking](#).

Related add-ons and applications

XenServer provides the core hypervisor features. To make your virtualization experience complete, you can add these related applications and services.

XenServer-specific add-ons:

- **XenCenter**

A Windows GUI client for VM management, implemented based on the management API. XenCenter provides a rich user experience to manage multiple XenServer hosts, resource pools, and the entire virtual infrastructure associated with them.

- **Workload Balancing (WLB)**

An appliance that balances your pool by relocating virtual machines onto the best possible hosts for their workload in a resource pool. For more information, see [Workload balancing \(/en-us/xenserver/9/wlb.html\)](/en-us/xenserver/9/wlb.html).

- **Conversion Manager**

A virtual appliance that enables users to convert existing VMware virtual machines into XenServer virtual machines, with comparable networking and storage connectivity. For more information, see [Conversion manager](#).

Related Citrix® products:

- **Citrix Licensing Server**

A Windows application that XenCenter contacts to request a license for the specified server and that collects telemetry about your XenServer environment.

- **License Activation Service**

The Citrix License Server reaches out to this service to verify what XenServer licenses your organization holds.

- **Citrix Provisioning™**

Provisioning Services that support PXE boot from common images. Used widely with Citrix Virtual Desktops and Citrix Virtual Apps. For more information, see [Provisioning](#).

- **Citrix Virtual Desktops™**

A Virtual Desktop Infrastructure (VDI) product specialized to Windows desktops. Citrix Virtual Desktops uses XAPI to manage XenServer in a multi-host pool configuration. For more information, see [Citrix Virtual Apps and Desktops](#).

Technical FAQs

June 9, 2026

Hardware

What are the minimum system requirements for running XenServer®?

For the minimum system requirements for this release, see [System requirements](#).

Do I need a system with a 64-bit x86 processor to run XenServer?

Yes. Either an Intel VT or AMD-V 64-bit x86-based system with one or more CPUs is required to run all [supported guest operating systems](#).

For more information about host system requirements, see [System requirements](#).

Do I need a system with hardware virtualization support?

You need a 64-bit x86 processor-based system that supports either Intel VT or AMD-V hardware virtualization technology in the processor and system firmware.

What systems are certified to run XenServer?

For a complete list of XenServer certified systems, see the [Hardware Compatibility List \(HCL\)](#).

Does XenServer support AMD Rapid Virtualization Indexing and Intel Extended Page Tables?

Yes. XenServer supports AMD Rapid Virtualization Indexing and Intel Extended Page Tables. Rapid Virtualization Indexing provides an implementation of nested tables technology used to further enhance the performance of the Xen® hypervisor. Extended Page Tables provide an implementation of hardware-assisted paging used to further enhance the performance of the Xen hypervisor.

Can XenServer run on a notebook or desktop-class systems?

XenServer runs on many notebook or desktop-class systems that conform to the minimum CPU requirements. However, XenServer only supports systems that have been certified and listed on the [Hardware Compatibility List \(HCL\)](#).

You can choose to run on unsupported systems for demonstration and testing purposes. However, some features, such as power management capabilities, do not work.

Can XenServer be installed on SD or USB cards?

No. XenServer does not support using SD cards or USB cards for your XenServer installation.

We only support hardware that has been certified and listed on the [Hardware Compatibility List \(HCL\)](#).

Do you support XenServer on end-of-life hardware?

Hardware that was certified and supported with a specific version of XenServer can reach EoL with the hardware vendor before the corresponding version of XenServer reaches its EoL.

We recommend that you do not install XenServer on EoL hardware. Instead source hardware that is supported by its vendor and XenServer. However, if the hardware was certified with your version of XenServer before its EoL date, we will attempt to investigate any reported issues with XenServer. While investigating a reported issue, if we determine that the issue directly relates to the EoL hardware, we will not investigate the issue further. This policy applies to constituent hardware components as well as to the overall system.

Ensure that you track the EoL dates of your hardware vendors and plan ahead to upgrade your hardware to later versions that are also certified and supported with the required versions of XenServer.

For more information, see the [Hardware Compatibility List \(HCL\)](#).

Product limits

For a complete list of XenServer supported limits, see [Configuration Limits](#).

Resource sharing

How are processing resources split between VMs?

XenServer splits processing resources between vCPUs using a fair-share balancing algorithm. This algorithm ensures that all VMs get their share of the processing resources of the system.

How does XenServer choose which physical processors it allocates to the VM?

XenServer doesn't statically allocate physical processors to any specific VM. Instead, XenServer dynamically allocates, depending on load, any available logical processors to the VM. This dynamic allocation ensures that processor cycles are used efficiently because the VM can run wherever there is spare capacity.

How are disk I/O resources split between the VMs?

XenServer uses a fair-share resource split for disk I/O resources between VMs. You can also provide a VM higher or lower priority access to disk I/O resources.

How are network I/O resources split between the VMs?

XenServer uses a fair-share resource split for network I/O resources between the VMs. You can also control the rate of outgoing data by using the Open vSwitch. For more information, see [Control the rate of outgoing data \(QoS\)](#).

Guest operating systems

Can XenServer run 32-bit operating systems as guests?

Yes. For more information, see [Supported guest operating systems](#).

Can XenServer run 64-bit operating systems as guests?

Yes. For more information, see [Supported guest operating systems](#).

Which versions of Microsoft Windows can run as guests on XenServer?

For a list of supported Windows guest operating systems, see [Supported guest operating systems](#).

Which versions of Linux can run as guests on XenServer?

For a list of supported Linux guest operating systems, see [Supported guest operating systems](#).

Can I run different versions of the supported operating systems or other unlisted operating systems?

We only support operating systems (OS) under OS vendor support. Although unsupported operating systems might continue to function, we might ask you to upgrade to a supported OS service pack before we can investigate any issues.

Applicable drivers might not be available for OS versions that are unsupported. Without the drivers, these OS versions do not function with optimized performance.

It's often possible to install other distributions of Linux. However, XenServer can only support the operating systems listed in [Supported guest operating systems](#). We might ask you to switch to a supported OS before issues we can investigate any issues.

Does XenServer support FreeBSD, NetBSD, or any other BSD variants as a guest operating system?

XenServer doesn't support any BSD-based guest operating systems for general-purpose virtualization deployments. However, FreeBSD VMs running on XenServer have been certified for use in specific NetScaler® products.

What are the XenServer VM Tools?

The XenServer VM Tools are software packages for Windows and Linux guest operating systems. For Windows operating systems, the XenServer VM Tools for Windows include high-performance I/O drivers (PV drivers) and the Management Agent.

For Linux operating systems, the XenServer VM Tools for Linux include a Guest Agent that provides additional information about the VM to the XenServer host.

For more information, see [XenServer VM Tools](#).

Docker

Can I run Docker containers on my Linux VMs?

Yes. Docker is supported on Linux VMs that are hosted on XenServer.

Can I run Docker containers on my Windows VMs?

No. You cannot run Docker containers on a Windows VM that is hosted on XenServer. This restriction is because XenServer does not support nested virtualization for Windows VMs.

XenCenter®

For more information, see [XenCenter](#).

Do I have to run XenCenter on a Windows computer?

Yes. The XenCenter management console runs on a Windows operating system. For information about the system requirements, see [System requirements](#)

If you don't want to run Windows, you can manage your XenServer hosts and pools by using the xe CLI or by using [xsconsole](#), a system configuration console.

Can I log on to XenCenter using my Active Directory user accounts?

Yes. You can set up XenCenter login requests to use Active Directory on all editions of XenServer.

For more information, see [Manage users](#).

Can I restrict access of certain functions within XenCenter to certain users?

Yes. The Role Based Access Control feature combined with Active Directory authentication can restrict access for users in XenCenter.

For more information, see [Manage users](#).

Can I use a single XenCenter console to connect to multiple XenServer hosts?

Yes. You can use a single XenCenter console to connect to multiple XenServer host systems.

Can I use XenCenter to connect to multiple hosts running different versions of XenServer?

Depending on the version of XenServer - yes. XenCenter is backwards-compatible with Citrix Hypervisor 8.0 and later versions. However, note that only XenServer 8.4 and later versions are receiving full support.

Can I use XenCenter to connect to multiple resource pools?

Yes. You can connect to multiple resource pools from a single XenCenter console.

How can I gain access to the console of a Linux VM?

The **Console** tab in XenCenter provides access to the text-based and graphical consoles of VMs running Linux operating systems. Before you can connect with the graphical console of a Linux VM, install and configure a VNC server and an X display manager on the VM.

XenCenter also enables you to connect to Linux VMs over SSH by using the **Open SSH Console** option on the **Console** tab of the VM.

How can I gain access to the console of a Windows VM?

XenCenter provides access to the emulated graphics for a Windows VM. If XenCenter detects remote desktop capabilities on the VM, XenCenter provides a quick connect button to launch a built-in RDP client that connects to the VM. Or, you can connect directly to your guests by using external remote desktop software.

Command line interface (CLI)

For more information, see [Command-line interface](#).

Does XenServer include a CLI?

Yes. All editions of XenServer include a full command line interface (CLI) –known as **xe**.

Can I access the xe CLI directly on the host?

Yes. You can access the CLI by connecting a screen and keyboard directly to the host, or through a terminal emulator connected to the serial port of the host.

Can I access the xe CLI from a remote system?

Yes. XenServer ships the xe CLI, which can be installed on Windows and 64-bit Linux machines to control XenServer remotely. You can also use XenCenter to access the console of the host from the Console tab.

Can I use the xe CLI using my Active Directory user accounts?

Yes. You can log in using Active Directory on all editions of XenServer.

Can I restrict access the use of certain CLI commands to certain users?

Yes. You can restrict user access on the xe CLI.

VMs

For more information, see [Manage virtual machines](#).

Can VMs created with VMware or Hyper-V run on XenServer?

Yes. You can export and import VMs using the industry-standard OVF format.

You can also convert VMs in batches using the Conversion Manager. Third-party tools are also available.

For more information, see [Conversion Manager](#).

What types of installation media can I use to install a guest operating system?

You can install a guest operating system by using:

- A CD in the CD-ROM drive of the host
- A virtual CD-ROM drive
- Placing ISO images on to a shared network drive
- Network installation, if supported by the specific guest.

For more information, see [Manage Virtual Machines](#).

Can I make a clone of an existing VM?

Yes. Any VM created on XenServer can be cloned or converted into a VM template. A VM template can then be used to create more VMs.

Can VMs be exported from one version of XenServer and moved to another?

Yes. VMs exported from older versions of XenServer can be imported to a newer version.

Can I convert a VM from the open-source version of Xen to XenServer?

No.

Does XenServer provide disk snapshot capabilities for VMs?

Yes. XenServer supports using snapshots in all editions. For more information, see [VM Snapshots](#).

Storage

For more information, see [Storage](#).

What types of local storage can be used with XenServer?

XenServer supports local storage such as SATA, SAS, and NVMe.

What type of SAN/NAS storage can be used with XenServer?

XenServer supports Fibre Channel, FCoE, Hardware-based iSCSI (HBA), iSCSI, NFS, and SMB storage repositories.

For more information, see [Storage](#) and the [Hardware Compatibility List](#).

Does XenServer support software-based iSCSI?

Yes. XenServer includes a built-in software-based iSCSI initiator (open-iSCSI).

What version of NFS is required for remote storage use?

XenServer requires NFSv3 or NFSv4 over TCP for remote storage use. XenServer currently does not support NFS over User Datagram Protocol (UDP).

Can I use software-based NFS running on a general-purpose server for remote shared storage?

Yes. Although we recommend using a dedicated NAS device with NFSv3 or NFSv4 with high-speed non-volatile caching to achieve acceptable levels of I/O performance.

Can I boot a XenServer host system from an iSCSI, Fibre Channel or FCoE SAN?

Yes. XenServer supports Boot from SAN using Fibre Channel, FCoE, or iSCSI HBAs.

Does XenServer support Multipath I/O (MPIO) for storage connections?

Yes. We recommend using multipath for resilient storage connections.

Does XenServer support a software-based RAID implementation?

No. XenServer doesn't support software RAID.

Does XenServer support HostRAID or FakeRAID solutions?

No. XenServer doesn't support proprietary RAID-like solutions, such as HostRAID or FakeRAID.

Does XenServer support thin cloning of existing VMs?

Yes. Thin cloning is available on local disks formatted as EXT3/EXT4, in addition to NFS and SMB storage repositories.

Does XenServer support Distributed Replicated Block Device (DRBD) storage?

No. XenServer doesn't support DRBD.

Does XenServer support ATA over Ethernet?

No. XenServer doesn't support ATA over Ethernet-based storage.

Networking

For more information, see [Networking](#)

Can I create private networks that isolate groups of VMs?

Yes. You can create a private network on a single host for resident VMs.

Does XenServer support multiple physical network connections?

Yes. You can connect to or associate multiple physical networks that attach to different network interfaces on the physical host system.

Can VMs connect to multiple networks?

Yes. VMs can connect to any network available to the host.

Does XenServer support IPv6?

VMs hosted on XenServer can use any combination of IPv4 and IPv6 configured addresses. You can also use the PVS-Accelerator with Provisioning targets in a site that stream over a mix of IPv4 and IPv6-based protocols. (IPv6 requires Citrix Provisioning™ 2402 CU1, Citrix Provisioning 2407, or later.)

However, XenServer doesn't support the use of IPv6 in its Control Domain (dom0). You can't use IPv6 for the host management network or the storage network. IPv4 must be available for the XenServer host to use.

Does XenServer support VLANs on a physical network interface?

Yes. XenServer supports assigning VM networks to specified VLANs.

Do XenServer virtual networks pass all network traffic to all VMs?

By default, XenServer network interfaces are non-promiscuous and a VM can only see traffic for that VM and broadcast traffic.

Does XenServer support bonding or teaming of physical network interfaces?

Yes. XenServer supports physical network interface bonding for failover and link aggregation with optional LACP support. For more information, see [Networking](#).

Memory**How much memory is consumed by running XenServer?**

Three components contribute to the memory footprint of a XenServer host.

1. The Xen hypervisor
2. The control domain on the host (dom0)
3. The XenServer Crash Kernel

The amount of memory required to run dom0 is adjusted automatically. By default, XenServer allocates 1 GiB plus 5% of the total physical memory to the control domain, up to an initial maximum of 8 GiB.

Note:

The amount of memory allocated to the Control Domain can be increased beyond the default amount.

In XenCenter, the **Xen** field in the **Memory** tab reports the memory used by the Control Domain, by the Xen hypervisor itself, and by the XenServer Crash Kernel. The amount of memory used by the hypervisor is larger for hosts with more memory.

For more information, see [Memory usage](#)

Does XenServer optimize VM memory usage?

Yes. XenServer uses Dynamic Memory Control (DMC) to automatically adjust the memory of running VMs. These adjustments keep the amount of memory allocated to each VM between specified minimum and maximum memory values, guaranteeing performance and permitting greater VM density.

For more information, see [VM memory](#).

Resource pools

For more information, see [Hosts and resource pools](#).

What is a resource pool?

A resource pool is a set of XenServer hosts managed as a unit. Typically, a resource pool shares some amount of networked storage to allow VMs to be rapidly migrated from one host to another within the pool.

Does XenServer require a dedicated host to manage a resource pool?

No. A single host in the pool must be specified as the pool coordinator. The pool coordinator controls all administrative activities required on the pool. This design means that there is no external single point of failure. If the pool coordinator fails, other hosts in the pool continue to operate, and the resident VMs continue to run as normal. If the pool coordinator cannot come back online, XenServer promotes one of the other hosts in the pool to coordinator to regain control of the pool.

This process is automated with the High Availability feature. For more information, see [High availability](#).

Where is the configuration data for a resource pool stored?

A copy of the configuration data is stored on every host in the resource pool. If the current pool coordinator fails, this data enables any host in the resource pool to become the new pool coordinator.

What types of configurations can be made at the resource pool level?

Shared remote storage and networking configurations can be made at the resource pool level. When a configuration is shared on the resource pool, the coordinator system automatically propagates configuration changes to all the member systems.

Are new host systems added to a resource pool automatically configured with shared settings?

Yes. Any new host systems added to a resource pool automatically receive the same configurations for shared storage and network settings.

Can I use different types of CPUs in the same XenServer resource pool?

Yes. We recommend that the same CPU type is used throughout the pool (homogeneous resource pool). However, it is possible for hosts with different CPU types to join a pool (heterogeneous), provided the CPUs are from the same vendor.

For more information, see [Hosts and resource pools](#).

For updated information about the support for feature masking for specific CPU types, see [Hardware Compatibility List](#).

Live Migration (formerly XenMotion)

For more information, see [Migrate VMs](#).

Can I move a running VM from one host to another?

With live migration you can move running VMs when hosts share storage (in a pool).

Also, storage live migration allows migration between hosts that do not share storage. VMs can be migrated within or across pools.

High availability

For more information, see [High availability](#).

Does XenServer offer high availability features?

Yes. If high availability is enabled, XenServer continually monitors the health of the hosts in a pool. If high availability detects that a host is impaired, the host is automatically shut down. This action allows for VMs to be restarted safely on an alternative healthy host.

Does XenServer high availability support local storage?

No. If you want to use high availability, shared storage is required. This shared storage enables VMs to be relocated if a host fails. However, high availability allows VMs that are stored on local storage to be marked for automatic restart when the host recovers after a reboot.

Can I use high availability to automatically sequence the restart of recovered VMs?

Yes. High availability configuration allows you to define the order that VMs are started. This capability enables VMs that depend on one another to be sequenced automatically.

Performance metrics

Do the XenServer management tools collect performance information?

Yes. XenServer provides detailed monitoring of performance metrics. These metrics include CPU, memory, disk, network, C-state/P-state information, and storage. Where appropriate, these metrics are available on a per-host and a per-VM basis. Performance metrics are available directly (exposed as Round Robin Databases), or can be accessed and viewed graphically in XenCenter or other third-party applications. For more information, see [Monitor and manage your deployment](#).

How are XenServer performance metrics gathered?

Data for the XenServer performance metrics are collected from various sources. These sources include the Xen hypervisor, Dom0, standard Linux interfaces, and standard Windows interfaces such as WMI.

Does XenCenter display performance metrics in real time?

Yes. XenCenter displays real-time performance metrics on the **Performance** tab for each running VM and for the XenServer host. You can customize the metrics that are displayed.

Does XenCenter store and display historic performance metrics?

Yes. XenServer keeps performance metrics from the last year (with decreasing granularity). XenCenter provides a visualization of these metrics in real-time graphical displays.

Installation

For more information, see [Install](#).

Does XenServer install on top of systems that are already running an existing operating system?

No. XenServer installs directly on bare-metal hardware, avoiding the complexity, overhead, and performance bottlenecks of an underlying operating system.

Can I upgrade an existing XenServer installation to a newer version?

Yes. If you are running a supported version of XenServer you can upgrade to a newer version of XenServer instead of doing a fresh installation. For more information, see [Upgrade](#).

Can I upgrade from an out-of-support version of Citrix Hypervisor or XenServer to this version?

If your existing version of Citrix Hypervisor or XenServer is no longer in support, you cannot upgrade or update to the latest version of XenServer.

For more information, see [Upgrade](#).

How much local storage does XenServer require for installation on the physical host system?

XenServer requires a minimum of 46 GB of local storage on the physical host system.

Can I use PXE to do a network installation of XenServer on the host system?

Yes. You can install XenServer on the host system by using PXE. You can also automatically install XenServer using PXE by creating a pre-configured answer file.

Does the Xen hypervisor run on Linux?

No. Xen is a Type 1 hypervisor that runs directly on the host hardware (“bare metal”). After the hypervisor loads, it starts the privileged management domain –the control domain (dom0), which contains a minimal Linux environment.

Where does XenServer get its device driver support?

XenServer uses the device drivers available from the Linux kernel. As a result, XenServer runs on a wide variety of hardware and storage devices. However, we recommend that you use certified device drivers.

For more information, see the [Hardware Compatibility List](#).

Licensing

For information about XenServer licensing, see [Licensing](#).

Lifecycle

For information about XenServer release lifecycle, see [XenServer Lifecycle](#).

Will XenServer 9 be a Long-Term Service Release (LTSR)?

XenServer 9 does not use the Current Release (CR) and Long Term Service Release (LTSR) strategy, as used on Citrix Virtual Apps and Desktops.

Instead, XenServer 9 uses an update stream model, as used on XenServer 8.4. For more information, see [Update your XenServer hosts](#).

XenServer 9 will be supported for 5 years from the General Availability date.

For more information about support dates, see [Product matrix](#).

Technical Support

For more information about support, see [Support](#).

Does XenServer provide direct technical support for XenServer?

Yes. For more information, visit [the XenServer support pages](#).

Do I have to purchase a XenServer technical support contract at the same time as I get a XenServer license?

No. A technical support contract is included with your license. For information about the level of support we provide for licensed customers, visit [the XenServer support pages](#).

What level of support does my license entitle me to?

When you get a XenServer per-socket license, you also get the benefit of our Technical Support services. For more information about levels of support, see <https://xenserver.com/support>.

Can I get support if my hosts are running in Trial Edition?

If you are a Trial Edition user, you are not eligible for support. However, we value your feedback: [Provide feedback](#).

Are there alternative channels for getting support for XenServer?

Yes. There are several alternative channels for getting technical support for XenServer. You can also use [Citrix Knowledge Center](#) or contract with authorized XenServer partners who offer technical support services.

Does XenServer provide technical support for the open-source Xen project?

No. XenServer doesn't provide technical support for the open-source Xen project. For more information, visit <http://www.xen.org/>.

What happens if an open-source component included within a XenServer release reaches EOL?

We support XenServer in accordance with the dates specified in the XenServer Product Matrix (<https://www.xenserver.com/support>). This includes open-source and 3rd-party components to such extent as they are included and used in our product.

What happens if a security scanner identifies an issue in the product?

We work hard to keep our product secure and endeavour to remediate any exploitable issues. We welcome reports of potential vulnerabilities from customers and security researchers.

Although we prioritize investigating and addressing issues that have a reasonable basis for exploitability over unsubstantiated reports, many security scanners do not attempt to determine if an issue is actually exploitable before reporting it as a vulnerability. If you are concerned about an issue reported by a security scanner, we appreciate any supporting evidence of potential exploitability.

Customers who have questions about reports from security scanners should contact their usual support channels. If you have grounds to believe that an issue is exploitable, you can report this to secure@cloud.com. See [Vulnerability Response](#) for details.

Can I open a technical support incident with XenServer if I'm experiencing a non-technical issue?

No. Raise any non-technical issues through Citrix® Customer Service. For example, issues to do with software maintenance, licensing, administrative support, and order confirmation.

Licensing overview

June 30, 2026

Important:

If you are using XenServer to run your Citrix Virtual Apps and Desktops workloads, you must have a Premium Edition license. For more information about getting a XenServer license, see the [XenServer website](#).

XenServer is an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more](#).

XenServer 9 is available in the following editions:

- **Premium Edition** is our premium offering, optimized for desktop, server, and cloud workloads. It offers the following features:
 - Dynamic Workload Balancing
 - GPU virtualization with NVIDIA vGPU
 - Thin provisioning for shared block storage devices
 - Support for SMB storage
 - Direct Inspect APIs
 - Export pool resource data
 - In-memory read caching

- PVS-Accelerator
 - Enablement for Citrix Virtual Desktops™ tablet mode
 - Changed block tracking
 - IGMP snooping
 - USB pass-through
 - SR-IOV network support
 - Monitor host and dom0 resources with NRPE
 - Monitor host and dom0 resources with SNMP
- **Trial Edition** showcases our Premium Edition feature set on a restricted size pool (maximum three hosts) for up to 90 days. If your pool has fewer than three hosts, you can add hosts until the pool reaches three. If your pool already has more than three hosts, you can keep the pool as it is without further expansion or host replacement. Trial Edition doesn't allow rolling pool upgrade through XenCenter. After the 90-day trial period expires, you cannot start any VMs on your Trial Edition host or pool.

For more information, see <https://www.xenserver.com/editions> and the [Feature Matrix](#).

	Premium Edition	Trial Edition
Premium feature set	Yes	Yes
Rolling pool upgrade through XenCenter	Yes	
Continuous functional and security updates	Yes	Yes
Maximum pool size	32 (16 with GFS2 SRs)	3
Other restrictions		Expires after 90 days

XenServer is licensed on a *per-socket* basis. Allocation of licenses is managed centrally and enforced by a standalone Citrix License Server in the environment based on information provided by the License Activation Service in Citrix Cloud. For more information about XenServer licensing, see the [Licensing FAQ](#).

Licensing your hosts and pools

XenServer uses the same licensing process as some Citrix products. To use XenServer Premium Edition, you require a valid license to be installed on a Citrix License Server and assigned to your XenServer host. This process is covered in detail in the [Licensing guide for XenServer](#).

Note:

XenServer does not currently support licensing hosted on Citrix Cloud. An on-premises Citrix® License Server is required.

To license your hosts, you need the following items:

- A Citrix Cloud account with the OrgID that is associated with your XenServer entitlements
- A Citrix License Server
- A XenServer host
- XenCenter

1. Install or upgrade to Citrix License Server version 11.17.2 build 53100 or later

You can download the Citrix License Server for Windows from the [Citrix Licensing downloads page](#). Ensure that you install or upgrade to version 11.17.2 build 53100 or later.

Install the Citrix License Server on a Windows system according to the instructions in the [Citrix Licensing documentation](#).

The Windows system that you install your Citrix License Server on can be a Windows VM hosted in your XenServer pool.

XenServer operates with a ‘grace’ license until the License Server can boot. This behavior means, after you have licensed the XenServer hosts in your pool, if you reboot the host that has the Citrix License Server running on it, a grace period is applied until the License Server is restarted.

2. Onboard with Citrix Cloud

Ensure that you use the OrgID associated with your XenServer entitlements. For more information, see the [Citrix Cloud product documentation](#).

3. Register your License Server with Citrix Cloud

Register your on-premises Citrix License Server by using short-code activation through Citrix Cloud. After registration, your XenServer entitlement can be found in the Citrix Cloud console. This entitlement enables you to activate XenServer by using the cloud. For more information, see the [Citrix Licensing product documentation](#) and [Citrix Cloud product documentation](#).

4. Ensure that your Citrix License Server and XenServer can communicate

Ensure that XenServer connects to the Citrix License Server on the specified IP and port.

By default, the Citrix License Server uses port 8083 for communication.

For information about the ports and protocols used for communication between these products, see [Communication ports used by XenServer product components](#).

Note:

Ensure that the time is synchronized between XenServer and the Citrix License Server.

5. Set up certificate trust for your Citrix License Server

When certificate verification is enabled, XenServer verifies the identity certificate presented by the Citrix License Server before checking out or reapplying a license.

After you apply the update for this feature, licensing continues to work as usual. If XenCenter cannot verify the License Server certificate when you assign a license, XenCenter displays the certificate and you can choose whether to trust it.

We recommend that you use a certificate signed by a trusted certificate authority (CA) and install the root CA certificate in the pool:

```
1 xe pool-install-trusted-certificate uuid=<pool_uuid> purpose=licensing  
   ca=true filename=<root-ca-for-license-server>.pem
```

If the License Server uses a self-signed certificate, or if you want to trust only that server certificate, install the License Server certificate as a pinned trusted certificate:

```
1 xe pool-install-trusted-certificate uuid=<pool_uuid> purpose=licensing  
   ca=false filename=<license-server-certificate>.pem
```

To view the trusted certificates for licensing, run the following commands:

```
1 xe certificate-list name="" purpose=licensing type=pinned  
2 xe certificate-list name="" purpose=licensing type=ca
```

If you upgraded your pool from an earlier version of XenServer, certificate verification is not enabled automatically. For more information, see [Certificate verification](#).

6. Apply the licenses to hosts in your resource pool

You can allocate licenses held in the License Activation Service to your XenServer hosts and pools by using XenCenter or the xe CLI.

Apply a license to all the hosts using XenCenter A: Follow this procedure to apply a license:

1. On the **Tools** menu, click **License Manager**.
2. Select the pool or hosts that you want to license, and then click **Assign License**.
3. In the **Apply License** dialog, specify the **Edition** type to assign to the host.
4. Type the host name or IP address of the Citrix License Server.
5. Click **OK**.

Apply a license to hosts or pools by using the xe CLI To apply a license to a single host, run the `host-apply-edition` command:

```
1  xe host-apply-edition edition=cloud \
2
3      license-server-address=<license_server_address> host-uuid=<
4          uuid_of_host> \
5      license-server-port=<license_server_port>
```

To apply a license to all hosts in a pool, run the `pool-apply-edition` command:

```
1  xe pool-apply-edition edition=cloud \
2
3      license-server-address=<license_server_address> pool-uuid=<
4          uuid_of_pool> \
5      license-server-port=<license_server_port>
```

Upgrade from file-based licensing to License Activation Service (LAS)

XenServer® has transitioned from file-based licensing to the License Activation Service (LAS), which provides centralized activation and streamlined license management.

1. Install or upgrade Citrix License Server to version **11.17.2 build 53100 or later**.
2. Register your License Server with **Citrix Cloud™** and ensure XenServer entitlement is included in your Customer Account of Citrix Cloud that was used in onboarding License Server.
 - a) Open the Citrix License Server web console, go to Dashboard → License Activation Service, and check whether **XenServer Premium Edition** appears in the product list.
3. Install **XenCenter version 2025.3 or later**.
4. Ensure that network communication is working and that time is synchronized between XenServer and the Citrix License Server. By default, XenServer connects to the Citrix License Server's IP address via port 8083.

5. Ensure that your XenServer pool/host trusts the certificate of the Citrix License Server.
6. Open **Tools** → **License Manager** in XenCenter.
7. Click **Assign License**.
8. Enter the **URL/IP address of the license server**, then click **OK**.

Auto-switch to License Activation Service (LAS)

XenServer can automatically switch a host from file-based (v6) or trial licensing to License Activation Service (LAS). If the prerequisites are met, the switch will occur automatically without requiring user configuration.

During a Toolstack restart, and periodically once a day, XenServer checks whether the host meets the prerequisites for LAS. If all prerequisites are met, the host is licensed using LAS. If any prerequisite is not met, the host remains licensed in file-based (v6) mode and XenServer raises an alert that explains why the switch did not occur. After you address the issue, the switch should occur when the next check is performed (for example on the following day).

In pooled environments, hosts might temporarily use different licensing modes if some hosts switch to LAS while others do not.

Note:

File-based licensing will no longer function after April 15, 2026.

Auto-switch a host to License Activation Service (LAS)

Use this task to understand how the automatic switch to LAS works and how to verify the host licensing state.

Before you begin, ensure the following requirements are met:

- The host is licensed in file-based (v6) mode or in trial mode with a configured license server address and port.
- The Citrix License Server version is 11.17.2 build 53100 or later.
- The License Server is registered with Citrix Cloud. The XenServer entitlement is included in the customer account used to onboard the License Server.

Procedure to force a check

1. Restart the Toolstack.
2. During the Toolstack restart, XenServer evaluates whether the host meets the prerequisites for LAS.

3. Observe the result:

- If all prerequisites are met, XenServer licenses the host using LAS.
- If one or more prerequisites are not met, the host remains licensed in file-based (v6) mode and XenServer raises an alert.

In XenCenter, confirm that the host licensing mode is License Activation Service (LAS). If the host remains licensed in file-based (v6) mode, review the alert details to identify why the automatic switch did not occur.

Troubleshooting If the automatic switch did not succeed, resolve the issues described in the alert. XenServer will retry the automatic switch the following day after the issues are resolved.


In pooled environments, if hosts are licensed using different licensing modes, manually switch the remaining hosts to LAS to restore a consistent licensing state across the pool.

The automatic switch to LAS might fail for the following reasons:

- The Citrix License Server version is earlier than 11.17.2 build 53100.
- The License Server is not registered with Citrix Cloud.
- The XenServer entitlement is not included in the customer account used to onboard the License Server.

Validate License Server Usage from Citrix Cloud In the Citrix Cloud Console, navigate to **Licensing -> Licensed deployments -> License Servers**.

Navigate to your License Server and select ‘...’-> **View Details**. Click the expand button (>) to see the list of activations.



Citrix XenServer			1 activation
ID	License Serve... ↑	Last successful connection	
ZL-XS-01	Connected	3/29/2026, 18:14:39 UTC	...

Showing 1-1 of 1 items Page 1 of 1 5 rows

ID: The table shows the XenServer Hosts activated with this License Server.

Click “...”-> **View Details** to see the granular details.

View details

Product & Version	CXS_ENT2_CCS:2024.0229
PoolUUID	8af0a411-4453-fe9e-3692-609b7b3249c3
HostUUID	cdeb1615-0742-4ca7-9b08-2eab6a26e8c5
Hostname	ZL-XS-01
LasUniqueEnvId	32a34eb88e690c990927b434fc337c083532662e11f9 34956ddd780194f8f6ff

Close

Pool UUID: The UUID of the XenServer pool that the XenServer host belongs to.

Host UUID: The UUID of the XenServer host

Hostname: The hostname of the XenServer host

Licensing FAQ

May 26, 2026

This article contains frequently asked questions about licensing your XenServer® hosts and pools.

- General questions
- Citrix Virtual Apps and Desktops™
- Citrix License Servers
- Licensing a XenServer pool
- More information

General questions

Q: Where can I get a XenServer license?

A: You can learn about getting a XenServer license at the [XenServer website](#).

Q: How do I apply a XenServer license?

A: XenServer requires a License Server. Connect your License Server to the License Activation Service in Citrix Cloud. The License Server can now view the licenses you have available for XenServer and assign them to your XenServer hosts.

When you assign a license to a XenServer host, XenServer contacts the specified Citrix License Server and requests a license for the specified hosts. If successful, a license is checked out and the License Manager displays information about the license the hosts are licensed under.

For more information, see [License Activation Service](#).

Q: How many licenses do I need to license my resource pool?

A: XenServer is licensed on a *per-CPU socket* basis. For a pool to be considered licensed, all XenServer hosts in the pool must be licensed. XenServer only counts populated CPU sockets.

You can use the Citrix License Server to view the number of available licenses displayed in the *License Administration Console Dashboard*.

Q: Do I need a per-socket license for sockets that are not populated?

A: No, only populated CPU sockets are counted toward the number of sockets to be licensed.

Q: Do I lose my virtual machine (VM) when my license expires?

A: No, you do not lose any VMs or their data.

Q: What happens if I have a licensed pool and the License Server becomes unavailable?

A. If your license has not expired and the License Server is unavailable, you receive a *grace period* of 15 days at the licensing level that was applied previously.

After the grace period, if your license is still unavailable, your pool changes to a Trial Edition pool and only those features included in Trial Edition are available for a 90-day period. Your pool remains the same size, but if it has three or more hosts, you cannot add any more hosts to it. After the 90-day trial period expires, you cannot start any VMs.

Q: What happens if I have a licensed pool, but XenServer cannot verify the identity of the License Server?

A. If your license has not expired and XenServer cannot verify the identity certificate presented by the License Server, you receive a *grace period* of 30 days at the licensing level that was applied previously.

After the grace period, if the License Server identity still cannot be verified, your pool changes to a Trial Edition pool and only those features included in Trial Edition are available for a 90-day period. Your pool remains the same size, but if it has three or more hosts, you cannot add any more hosts to it. After the 90-day trial period expires, you cannot start any VMs.

If the certificate check fails during automatic license reapply or after a toolstack restart, XenCenter shows the alert **License server certificate check failed**. To resolve the issue, reassign the License Server in XenCenter or install a trusted certificate by using the xe CLI.

If XenCenter shows a certificate that you recognize and trust, accept it to install the certificate in the pool as a pinned trusted certificate. If the certificate is unexpected, do not trust it. Check the License Server certificate configuration and investigate whether the connection might be intercepted.

You can also install trusted certificates by using the xe CLI:

- Install a root CA certificate for licensing when the License Server uses a CA-signed certificate:

```
1 xe pool-install-trusted-certificate uuid=<POOL_UUID> purpose=
   licensing ca=true filename=<root-ca-for-license-server>.pem
```

- Install a pinned certificate when the License Server uses a self-signed certificate, or when you want to trust only the presented server certificate:

```
1 xe pool-install-trusted-certificate uuid=<POOL_UUID> purpose=
   licensing ca=false filename=<license-server-certificate>.pem
```

- View the trusted certificates for licensing:

```
1 xe certificate-list name="" purpose=licensing type=pinned
2 xe certificate-list name="" purpose=licensing type=ca
```

- If you install a root CA certificate and the pool already has a pinned certificate for the same License Server, remove the pinned certificate so XenServer can use the CA certificate:

```
1 xe pool-uninstall-trusted-certificate uuid=<POOL_UUID>
   certificate-uuid=<CERTIFICATE_UUID>
```


Q: What happens if I have a licensed pool and the License Server cannot contact Citrix Cloud?

A: If your license has not expired and the License Server cannot contact Citrix Cloud, you receive a *grace period* of 30 days at the licensing level that was applied previously.

After the grace period, if your license is still unavailable, your pool changes to a Trial Edition pool and only those features included in Trial Edition are available for a 90-day period. Your pool remains the same size, but if it has three or more hosts, you cannot add any more hosts to it. After the 90-day trial period expires, you cannot start any VMs.

Q: How do I get a license to evaluate XenServer?

A: With Trial Edition, you can install XenServer without a license and trial all the Premium Edition features for 90 days in a reduced-size pool. For more information, see [XenServer editions](#).

After the 90-day trial period expires, you cannot start any VMs on your Trial Edition host or pool.

Q: Can I use XenServer 9 without a license?

A: Yes. XenServer 9 Trial Edition makes XenServer available without a license. The edition provides all of the Premium Edition features for 90 days in a reduced-size pool of up to three hosts. For more information, see [XenServer editions](#).

After the 90-day trial period expires, you cannot start any VMs on your Trial Edition host or pool.

Q: If I'm already using XenServer with Trial Edition without the trial period expiry, when does the trial period begin?

After you apply the updates that introduce the expiry period, the trial period begins on the updated hosts and pools. This period expires 90 days after applying the update.

Q: Can I start an existing shutdown VM after the 90-day limit?

VMs cannot be started on your Trial Edition hosts and pools after the 90-day trial period expires. This includes new VMs and VMs that were created before the trial period expires.

Q: Can I use Citrix Cloud to apply a license to XenServer?

To license XenServer, you require a License Server version 11.17.2 build 53100 or later. The License Server connects to the License Activation Service in Citrix Cloud and activates your XenServer license. For more information, see [License Activation Server](#).

Q: What level of support does my license entitle me to?

When you get a XenServer per-socket license, you also get the benefit of our Technical Support services. For more information about levels of support, see <https://xenserver.com/support>.

If you are a Trial Edition user, you are not eligible for support. However, we value your feedback: [Provide feedback](#).

For more information, see [Technical Support questions](#) and [Support](#)

Citrix Virtual Apps and Desktops

Q: What edition of XenServer do I need to run Citrix Virtual Apps and Desktops workloads on XenServer?

A: To run your Citrix Virtual Apps and Desktops workloads on XenServer, you must have a Premium Edition per-socket license for XenServer. Premium Edition delivers many advanced features that make XenServer a highly optimized hypervisor platform for your workload. For more information, see [XenServer editions](#).

For more information about getting a XenServer license, see the [XenServer website](#).

XenServer is now an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more](#).

Q: What are the constraints on the use of the XenServer advanced virtualization management capabilities delivered as part of Citrix Virtual Apps and Desktops?

A: To run Citrix Virtual Apps and Desktops workloads on XenServer, you need a Premium Edition license. Premium Edition delivers many advanced features that make XenServer a highly optimized hypervisor platform for your workload.

For information about the advanced capabilities that XenServer provides for Citrix Virtual Apps and Desktops, see [Using XenServer with Citrix products](#).

For more information about getting a XenServer license, see the [XenServer website](#). XenServer is now an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more](#).

Citrix License Servers

Q: Which License Servers can I use with XenServer?

A: Use the Citrix License Server software version 11.17.2 build 53100 or later on a server running Microsoft Windows. This version or later is required for licensing with the License Activation Service.

Q: How do I import my license onto the Citrix License Server?

A: For information on importing a license file, see the [Citrix Licensing documentation](#).

Q: Can I run the License Server on my XenServer pool?

A: Yes. You can install the Citrix License Server software on a Windows VM.

XenServer operates with a ‘grace’ license until the License Server is able to boot. This behavior means, after you have licensed the XenServer hosts in your pool, and you reboot the host that has the Citrix License Server running on it, a grace period is applied to that host until the License Server is restarted.

Q: Can I use the Windows version of the Citrix License Server with XenServer?

A: Yes.

Q: Can I install Licenses for Citrix products on the Citrix License Server software installed on Windows?

A: Yes, you can license Citrix products by using the Citrix License Server software installed on Windows. For more information, see [Licensing](#) on the [Citrix Product Documentation](#) website.

Licensing a XenServer pool

Q: How do I apply a license to all the hosts using XenCenter®?

A: Follow this procedure to apply a license:

1. On the **Tools** menu, click **License Manager**.
2. Select the Pool or Hosts you want to license, and then click **Assign License**.
3. In the **Apply License** dialog, specify the **Edition** type to assign to the host, and type the host name or IP address of the License Server.

Q: Can I apply a license without using XenCenter?

A: Yes, you can use the xe CLI. Run the `host-apply-edition` command. For example, enter the following to license a host:

```
1 xe host-apply-edition edition=premium-per-socket|standard-per-socket|  
  cloud \  
2  
3   license-server-address=<license_server_address> host-uuid=<  
    uuid_of_host> \  
4  
5   license-server-port=<license_server_port>
```

To license a pool, use the `pool-apply-edition` command. For example, enter the following to license a pool:

```
1 xe pool-apply-edition edition=premium-per-socket|standard-per-socket|  
  cloud \  
2  
3   license-server-address=<license_server_address> host-uuid=<  
    uuid_of_host> \  
4  
5   license-server-port=<license_server_port>
```

Q: How can I discover the license status of my hosts and pools?

A: XenCenter displays the license type of a host or pool.

To see the license type of a host or pool, select that host or pool in the tree view. XenCenter displays the license status in the title bar for that host or pool, after the host or pool name.

You can also go to the **General** tab of the host and find the license type in the **License Details** section.

To find the license type of a host by using the command line, run the following command in the console of a host in your pool:

```
1 xe host-license-view host\_uuid=<UUID> | grep sku\_marketing\_name
```

More information

- For more information about the XenServer 9 release, see [XenServer 9 Release Notes](#).
- To access XenServer 9 product documentation, see [XenServer 9 Product Documentation](#).
- For an overview of the XenServer product, see [Technical Overview](#).
- For support information, see <https://xenserver.com/support>.

Install

June 29, 2026

XenServer® installs directly on bare-metal hardware avoiding the complexity, overhead, and performance bottlenecks of an underlying operating system.

XenServer uses the device drivers available from the Linux kernel. As a result, XenServer can run on a wide variety of hardware and storage devices. However, ensure that you use certified device drivers. For more information, see the [Hardware Compatibility List \(HCL\)](#).

Important:

The XenServer host must be installed on a dedicated 64-bit x86 server. Do not install any other operating system in a dual-boot configuration with the XenServer host. This configuration is not supported.

This section is primarily aimed at system administrators who want to set up XenServer hosts on physical servers. It contains procedures to guide you through the installation or upgrade process. It also contains information about troubleshooting problems that might occur during installation and points you to extra resources.

Before you start

Depending on your environment, the installation method to use to get your hosts and pools to the latest version of XenServer 9 differs.

- If you already have XenServer 8.4 installed on your hosts and pools, [Upgrade from XenServer 8.4](#).
- If you already installed XenServer 9, you cannot upgrade or update from the installation ISOs. Instead, apply the latest level of frequent updates through XenCenter. For more information, see [Update XenServer 9](#).
- If you are already running any other version of XenServer or Citrix Hypervisor on your hosts and pools, upgrading from these versions is not supported. Do a fresh installation of XenServer 9.
- If you are installing XenServer for the first time on your hosts and pools, do a fresh installation of XenServer 9.

Installation methods

XenServer 9 can be installed in one of the following ways:

Fresh installation If you are creating a fresh installation of XenServer 9:

- Use the **XenServer 9 Installation ISO** file. You can download this file from the [XenServer downloads page](#).
- Review the information in the [Hardware Compatibility List \(HCL\)](#) and your hardware vendor documentation to ensure that your hardware is supported and is installed with a firmware version that is compatible with the drivers included in the XenServer 9 installation ISO. For more information, see [Drivers](#).
- Review the information in [System Requirements](#), [Licensing XenServer](#), and [Installing XenServer and XenCenter](#) before installing XenServer.

Upgrade If you are upgrading from XenServer 8.4 to XenServer 9:

- Use the **XenServer 9 Installation ISO** file. You can download this file from the [XenServer downloads page](#).
- Review the information in the [Hardware Compatibility List \(HCL\)](#) and your hardware vendor documentation to ensure that your hardware is supported and is installed with a firmware version that is compatible with the drivers included in the XenServer 9 installation ISO. For more information, see [Drivers](#).
- Review the information in [System Requirements](#), [Licensing XenServer](#), and [Upgrading from an existing version](#) before upgrading XenServer.

The installer presents the option to upgrade when it detects a previously installed version of XenServer. The upgrade process follows the first-time installation process, but several setup steps are bypassed. The existing settings are retained, including networking configuration, system time and so on.

You cannot upgrade directly from out-of-support versions of XenServer or Citrix Hypervisor to XenServer 9. Instead, perform a fresh installation.

Update Existing installations of XenServer 9 receive the latest updates through the frequent update mechanism. For more information, see [Update XenServer 9](#).

Supported boot modes

XenServer supports booting hosts using UEFI boot mode. For more information, see [Secure Boot for XenServer 9](#)

Install the XenServer host

This procedure takes you through doing a manual installation from local media. For information about other types of installation - such as network installation, unattended installation, or boot from SAN, see [Other installation scenarios](#).

Tip:

Throughout the installation, quickly advance to the next screen by pressing **F12**. Use the **Tab** key to move between elements and **Space** or **Enter** to select. Press **F1** for general help.

To install a XenServer host:

1. Back up any data you want to preserve. Installing XenServer overwrites data on any hard drives that you select to use for the installation.
2. Boot the computer from the installation media:
 - To install XenServer host from a bootable USB:
 - a) Create a bootable USB from the XenServer installation ISO.
Ensure that the tool does not alter the contents of the ISO file.
If you use a tool that modifies the image, you might see a **GRUB** screen at boot.
If this occurs, try again using a different tool that does not modify the image.
 - On Linux, you can use the **dd** command to write the ISO to a USB. For example, `dd if=<path_to_source_iso> of=<path_to_destination_usb>`.
 - On Windows, you can use Rufus. Ensure that you select **Write in DD Image mode**. If this is not selected, Rufus can alter the contents of the ISO file and cause it not to boot.
 - b) Insert the bootable USB drive into the target system.
 - c) Restart the system.
 - d) Go into the boot menu.
 - e) Change the settings to boot the system from the USB.
(If necessary, see your hardware vendor documentation for information on changing the boot order)
 - To install XenServer host from a CD/DVD:
 - a) Burn the XenServer installation ISO file to a CD/DVD.
 - b) Insert the bootable CD/DVD into the CD/DVD drive on the target system.
 - c) Restart the system.

- d) Go into the boot menu.
- e) Change the settings to boot the system from the CD/DVD.
(If necessary, see your hardware vendor documentation for information on changing the boot order)
- To install XenServer host from virtual media:
 - a) Go to the virtual console of your system.
 - b) Insert the XenServer installation ISO file as virtual media.
 - c) Restart the system.
 - d) Go into the boot menu.
 - e) Change the settings to boot the system from the virtual media.
(If necessary, see your hardware vendor documentation for information on changing the boot order)
- To install XenServer from a network location, see [Network installation](#).

3. Following the initial boot messages and the **Welcome to XenServer** screen, select your key map (keyboard layout) for the installation.

Note:

If a System Hardware warning screen is displayed and hardware virtualization assist support is available on your system, see your hardware manufacturer for BIOS upgrades.

4. At the **Welcome to XenServer Setup** screen, select **OK** to proceed.
5. At the **Select Driver** screen, select device drivers required for your hardware, and then select **OK** to proceed.
6. Scroll through and read the XenServer End User Agreement (EUA). Select **Accept EUA** to proceed.
If you choose not to accept the EUA, you cannot continue with the installation.
7. Select the appropriate action from the list. This list always includes:

- **Perform clean installation:** Choose this option to continue with a fresh installation.

Depending on the state of your server, you might also see the following options:

- **Upgrade:** If the installer detects a previously installed version of XenServer, it offers the option to upgrade. For information about upgrading your XenServer host, see [Upgrading from an existing version](#).

- **Restore:** If the installer detects a previously created backup installation, it offers the option to restore XenServer from a backup.

Make your selection, and choose **OK** to proceed.

8. If you have multiple local hard disks, choose a primary disk for the installation. Select **OK**.
9. Choose which disks that you want to use for virtual machine storage. View information about a specific disk by pressing **F5**. Select **OK**.
10. If you selected disks that all have 512 byte blocks, you have the option to use thin provisioning to optimize the use of available storage. Select **Enable thin provisioning** to enable the local SR of the host to be used for the local caching of VM VDI's (the Intellicache feature). Citrix Virtual Desktops and DaaS users are recommended to select this option for local caching to work properly. For more information, see [Intellicache](#).

Note:

If you selected disks that are 4KB native, the virtual machine storage is automatically configured for a large disk block size.

Select **OK** to proceed.

11. Select your installation media source.
 - To install from a USB, CD, or virtual media, select **Local media**.
 - To install from the network, select **HTTP** or **FTP** or **NFS**.

Select **OK** to proceed.

12. If you select HTTP or FTP or NFS in the previous step, set up networking so that the installer can connect to the XenServer installation media files:
 - a) If the computer has multiple NICs, select one of them to be used to access the XenServer installation media files. Choose **OK** to proceed.
 - b) Choose **Automatic configuration (DHCP)** to configure the NIC using DHCP, or Static configuration to configure the NIC manually. If you choose **Static configuration**, enter details as appropriate.
 - c) Provide VLAN ID if you have your installation media present in a VLAN network.
 - d) If you choose **HTTP** or **FTP**, provide the URL for your HTTP or FTP repository, and a user name and password, if appropriate.

If you choose **NFS**, provide the server and path of your NFS share.
 - e) Choose **OK** to proceed.

For more information about setting up your installation media on NFS, FTP, or HTTP, see [Other installation scenarios](#).

13. Indicate if you want to verify the integrity of the installation media. If you select **Verify installation source**, the SHA-256 checksum of the packages is calculated and checked against the known value. Verification can take some time. Make your selection and choose **OK** to proceed.
14. Set and confirm a root password, which XenCenter® uses to connect to the XenServer host. You also use this password (with user name “root”) to log into **xsconsole**, the system configuration console.

Note:

XenServer root passwords must contain only printable ASCII characters.

15. Set up the primary management interface that XenCenter uses to connect to and manage the host.

If your computer has multiple NICs, select the NIC that you want to use for management. Choose **OK** to proceed.

16. Configure the management interface with the following options:
 - Choose **Automatic configuration (DHCP)** to configure the NIC using DHCP.
 - Choose **Static configuration** to configure the NIC manually. Provide an IP address, subnet mask, and gateway.
 - Select **Use VLAN** to have the management interface on a VLAN network. Provide the VLAN ID.

Note:

To be part of a pool, XenServer hosts must have static IP addresses or be DNS addressable. When using DHCP, ensure that a static DHCP reservation policy is in place.

17. Specify the host name and the DNS configuration.

a) In the **Hostname Configuration** section, select from the following options:

- Choose **Automatically set via DHCP** to have the DHCP server provide the host name along with the IP address.
- Choose **Manually specify** to define the host name yourself. Enter the host name for the server in the field provided.

Note:

If you manually specify the host name, enter a short host name and *not the fully qualified*

domain name (FQDN). Entering an FQDN can cause external authentication to fail, or the XenServer host might be added to Active Directory with a different name.

a) In the **DNS Configuration** section, select from the following options:

- Choose **Automatically set via DHCP** to get name service configuration using DHCP.
- Select **Manually specify** to define the DNS servers yourself. Enter the IP addresses of your primary (required), secondary (optional), and tertiary (optional) DNS servers in the fields provided.

Select **OK** to proceed.

18. Select your time zone by geographical area and city. You can type the first letter of the desired locale to jump to the first entry that begins with this letter. Choose **OK** to proceed.

19. Specify how you want the XenServer host to determine local time. XenServer offers the following options:

- **Using NTP:** Select this option to use the NTP protocol to set your server time. On the next screen, configure it in one of the following ways:
 - Select **NTP is configured by my DHCP server** to have your network provide the NTP server host name or IP address.
 - Manually enter at least one NTP server name or IP address.

Choose **OK** to proceed.

- **Manual time entry:** Select this option to set the date and time manually.

On the next screen, enter the current time in UTC.

Choose **OK** to proceed.

Note:

XenServer assumes that the time setting on the server is the current time in UTC.

Choose **OK** to proceed.

20. Select **Install XenServer**.

The installation process starts. This process might take some minutes.

21. The next screen asks if you want to install any supplemental packs.

- If you want to install any supplemental packs, choose **Yes**.

Important:

If you want to install the NVIDIA Virtual GPU Manager supplemental pack, we strongly recommend that you do not install this pack now. Instead, install it into your

XenServer host by using `xe` CLI or XenCneter. For more information, see [Install the NVIDIA vGPU Manager for XenServer](#).

- a) You are prompted to insert the supplemental pack. Eject the XenServer installation media, and insert the supplemental pack media.
 - b) Choose **OK**.
 - c) Select **Use media** to proceed with the installation.
 - d) Repeat for each pack to be installed.
- If you do not want to install a supplemental pack, choose **No**.

The installation process completes. This process can take a few minutes.

22. When prompted by the **Installation Complete** screen, eject the installation media (if installing from USB or CD).
23. Select **OK** to reboot the host.

After the host reboots, XenServer displays **xsconsole**, a system configuration console. To access a local shell from **xsconsole**, press **Alt+F3**; to return to **xsconsole**, press **Alt+F1**.

Note:

Make note of the IP address displayed. Use this IP address when you connect XenCenter to the XenServer host.

If you encounter issues during your installation, gather the logs as described in this article: [Installation logs](#).

Install XenCenter

XenCenter must be installed on a Windows machine that can connect to the XenServer host through your network. Ensure that .NET framework version 4.8 or above is installed on this system.

To install XenCenter:

1. Download the installer for the latest version of XenCenter from the [XenServer downloads page](#).
2. Launch the installer `.msi` file.
3. Follow the **Setup** wizard, which allows you to modify the default destination folder and then to install XenCenter.

For more information about using XenCenter, see the [XenCenter documentation](#).

Connect XenCenter to the XenServer host

To connect XenCenter to the XenServer host:

1. Launch XenCenter. The program opens to the **Home** tab.
2. Click the **Add New Server** icon.
3. Enter the IP address of the XenServer host in the **Server** field. Type the root user name and password that you set during XenServer installation. Click **Add**.
4. The first time you add a host, the **Save and Restore Connection State** dialog box appears. This dialog enables you to set your preferences for storing your host connection information and automatically restoring host connections.

If you later want to change your preferences, you can do so from the XenCenter main menu, select **Tools** and then **Options**. The **Options** dialog box opens. Select the **Save and Restore** tab and set your preferences. Click **OK** to save your changes.

License your XenServer hosts

Your newly installed XenServer hosts can run in Trial Edition without a license for up to 90 days. This edition restricts your pool size to 3 and doesn't allow rolling pool upgrade through XenCenter. After the 90-day trial period expires, you cannot start any VMs on your Trial Edition host or pool. For more information, see <http://www.xenserver.com/editions>.

For information about licensing your XenServer hosts, see [Licensing XenServer](#).

Other installation scenarios

June 22, 2026

In addition to a standard manual installation process, XenServer® provides the ability to perform various other types of installations, including the following:

- Network installations using PXE boot
- Unattended installations
- Setting up the host to boot from SAN
- Configuring host multipathing

If you encounter issues during your installation, gather the logs as described in this article: [Installation logs](#).

Supported boot modes

XenServer supports booting hosts using UEFI or UEFI Secure Boot mode. For more information on Secure Boot, see [Secure Boot for XenServer 9](#).

Network installation

If the server that you want to install on has a PXE boot-enabled Ethernet card, you can use this feature to do a network installation with PXE boot.

Using PXE boot to install from the network involves the following steps:

- Copy the installer files to a TFTP server and configure your TFTP and DHCP servers for PXE boot installation.
- Host your installation media on NFS, FTP, or HTTP. Only the installer files are accessed from the TFTP server. The XenServer files to be installed on the server are hosted on an NFS, FTP, or HTTP server. Alternatively, after you start the install through PXE boot, you can complete it from local media hosted on the target server.
- (Optional) Create an answer file for unattended installation. You can instead choose to do an attended installation and step through the installer manually.
- Start the installation process.

Note:

PXE boot is not supported over a tagged VLAN network. Ensure that the VLAN network you use for PXE boot is untagged.

Configure your TFTP and DHCP servers

Before you set up the XenServer installation media, configure your TFTP and DHCP servers. The following sections contain information on how to configure your TFTP server for PXE boot with UEFI. Consult your vendor documentation for general setup procedures.

Configure your TFTP server for PXE boot with UEFI Host the installer files on a TFTP server and configure your DHCP and TFTP servers to enable PXE booting with UEFI boot mode. This configuration is used to start the installation process.

1. In the TFTP root directory (for example, `/tftpboot`), create a directory called `EFI/xenserver`.

2. Copy the following files from the XenServer installation media to the new `EFI/xenserver` directory on the TFTP server:

- `grubx64.efi` from the `/EFI/xenserver` directory
- `install.img` from the root directory
- `mlinuz` from the `/boot` directory
- `xen.gz` from the `/boot` directory

3. Configure your DHCP server to provide `/EFI/xenserver/grubx64.efi` as the boot file.

4. In the `EFI/xenserver` directory on the TFTP server, create the `grub.cfg` file.

The content of this file depends on how you want to configure your PXE boot environment and on the values that are appropriate for your servers.

- **Example: Unattended install** This example configuration performs an unattended installation using the answer file at the URL specified:

```
1  menuentry "XenServer Install (serial)" {
2
3      multiboot2 /EFI/xenserver/xen.gz dom0_max_vcpus=1-16
4          dom0_mem=max:8192M com1=115200,8n1 console=com1,vga
5      module2 /EFI/xenserver/mlinuz console=hvc0 console=tty0
6          answerfile_device=<interface_name> answerfile=http://<
          ip_address>/<path_to_answer_file> install
7      module2 /EFI/xenserver/install.img
8  }
```

Note:

To specify which network adapter to use to retrieve the answer file, include the `answerfile_device=<interface_name>` or `answerfile_device=<MAC>` parameter and specify either the network interface name (for example, `eno12429np0`) or the MAC address of the device.

In environments where network adapters may be enumerated differently across hosts (for example, when using different hardware generations), you can specify the `map_netdev` boot parameter to explicitly control the logical ordering of network adapters based on MAC address.

The `map_netdev` parameter must be supplied as an installer boot parameter and is commonly used with unattended installations to ensure consistent network adapter ordering across hosts. This parameter supports mapping a specific MAC address to logical index (e.g., `map_netdev=0:<MAC>`).

To ensure a deterministic NIC order, repeat this parameter for each required adapter. For example, you can specify `map_netdev=0:00:11:22:AA:BB:CC` to force a

specific hardware mapping.

For more information about using an answer file, see [Create an answer file for unattended installation](#).

- **Example: Manual install** This example configuration starts an installation that boots from the TFTP server and requires manual responses:

```
1  menuentry "XenServer Install (serial)" {  
2  
3      multiboot2 /EFI/xenserver/xen.gz dom0_max_vcpus=1-16  
          dom0_mem=max:8192M com1=115200,8n1 console=com1,vga  
4      module2 /EFI/xenserver/vmlinuz console=hvc0 console=tty0  
5      module2 /EFI/xenserver/install.img  
6  }
```

Next step: Host your installation media on NFS, FTP, or HTTP. In addition to the TFTP and DHCP servers, you require an NFS, FTP, or HTTP server to house the XenServer files that are installed on your server.

Host your installation media on NFS, FTP, or HTTP

The TFTP server hosts the files needed to start the installer, but the files to be installed or to use for an upgrade are hosted on an NFS, FTP, or HTTP server.

You can also use files hosted on NFS, FTP, or HTTP to complete an installation or upgrade that has been started from local media on your server.

1. On the HTTP, FTP, or NFS server, create a directory from which the XenServer installation media can be exported through HTTP, FTP, or NFS.
2. If you are using IIS to host the installation media, assuming a virtual directory is used, complete the following set up:
 - a) Enable directory browsing for the virtual directory.
 - b) Ensure that double escaping is enabled for the virtual directory before extracting the installation ISO on it.
 - c) Configure the following required MIME types
 - i. . as an application/octet-stream
 - ii. .* as an application/octet-stream

The following [web.config](#) file demonstrates these values configured for a virtual directory


```
1 <?xml version="1.0" encoding="UTF-8"?>
2 <configuration>
3   <system.webServer>
4     <security>
5       <requestFiltering allowDoubleEscaping="true" />
6     </security>
7     <directoryBrowse enabled="true" showFlags="Date, Time,
      Size, Extension, LongDate" />
8     <staticContent>
9       <mimeTypeMap fileExtension="." mimeType="application/
      octet-stream" />
10      <mimeTypeMap fileExtension="*" mimeType="application/
      octet-stream" />
11    </staticContent>
12  </system.webServer>
13 </configuration>
```

3. Unpack the contents of the XenServer installation ISO. You can do this in one of the following ways:
 - By using an ISO decompression application, such as 7-zip.
 - By mounting the ISO and copying all of the contents.
4. Copy the entire unpacked contents of the XenServer installation ISO to the newly created directory on the HTTP, FTP, or NFS server. This directory is your installation repository.

Note:

When copying the XenServer installation media, ensure that you copy the file `.treeinfo` to the newly created directory.

Next step:

- If you are completing an unattended installation: Create an answer file for unattended installation.
- If you are using PXE boot to start a manual installation: Start the network installation.

Create an answer file for unattended installation

To perform installations in an unattended fashion, create an XML answer file.

Contain all nodes within a root node named *installation*. When constructing your answer file, refer to the Answer file reference.

Here is an example answer file:

```
1 <?xml version="1.0"?>
2   <installation srtype="ext">
```

```
3      <primary-disk>sda</primary-disk>
4      <guest-disk>sdb</guest-disk>
5      <guest-disk>sdh</guest-disk>
6      <keymap>us</keymap>
7      <root-password>mypassword</root-password>
8      <source type="url">http://pxehost.example.com/xenserver/</
      source>
9      <script stage="filesystem-populated" type="url">
10         http://pxehost.example.com/myscripts/post-install-script
11      </script>
12      <admin-interface name="<interface-name>" proto="dhcp" />
13      <timezone>Europe/London</timezone>
14  </installation>
```

Next step: Start the network installation.

Automated upgrades with an answer file You can also perform automated upgrades by changing the answer file appropriately.

1. Set the `mode` attribute of the `installation` element to `upgrade`.
2. Specify the disk on which the existing installation lives with the `existing-installation` element.
3. Leave the `primary-disk` and `guest-disk` elements unspecified.

For example:

```
1  <?xml version="1.0"?>
2  <installation mode="upgrade">
3      <existing-installation>sda</existing-installation>
4      <source type="url">http://pxehost.example.com/xenserver/</source>
5      <script stage="filesystem-populated" type="url">
6          http://pxehost.example.com/myscripts/post-install-script
7      </script>
8  </installation>
```

Answer file reference The following is a summary of the elements. All node values are text, unless otherwise stated. Required elements are indicated.

<installation> **Required?** Yes

Description: The root element that contains all the other elements.

Attributes:

`srtype`

The attribute `srtype` can have one of the following values: `lvm`, `ext`, or `xfs`:

- `lvm` - set the local storage type to LVM.
- `ext` - set the local storage type to EXT4. This enables local caching for Citrix Virtual Desktops to work properly. For more information, see [Storage](#).
- `xfs` - set the local storage type to XFS. This option also allows you to create local storage devices with 4 KB physical blocks without requiring a logical block size of 512 bytes.

To enable thin provisioning, you can specify the `srtype` attribute as `ext` or `xfs`. If you do not specify the `srtype` attribute, the default value for `srtype` is `lvm`. If you do not specify the `srtype` attribute but you configure a 4 KB native disk for local storage in your answer file, the default value is `xfs`.

Note:

You cannot use Local LVM or Local EXT3/EXT4 storage types with 4 KB physical blocks. If you attempt to specify `lvm` or `ext` for the `srtype` attribute whilst configuring 4 KB physical blocks, your answer file configuration is rejected as incompatible.

`mode`

To change the installation type to upgrade, specify a `mode` attribute with the value `upgrade`. If this attribute is not specified, the installer performs a fresh installation and overwrites any existing data on the server.

`<primary-disk>` **Required?** Yes

Note:

Deprecated for upgrade scenarios.

Description: The name of the storage device where the control domain is installed. This element is equivalent to the choice made on the *Select Primary Disk* step of the manual installation process.

Attributes: You can specify a `guest-storage` attribute with possible values `yes` and `no`. For example: `<primary-disk guest-storage="no">sda</primary-disk>`

The default value is `yes`. If you specify `no`, you can automate an installation scenario where no storage repository is created. In this case, specify no guest-disk keys.

`<guest-disk>` **Required?** No

Description: The name of a storage device to be used for storing guests. Use one of these elements for each extra disk.

Attributes: None

<ntp> Required? Yes

Description: Specifies the source for NTP servers. If the `<ntp>` element is not specified, the default shall be `manual` if `<ntp-server>` is specified, `dhcp` if using DHCP, otherwise `default`.

Attributes:

The attribute `source` can have one of the following values: `dhcp`, `default`, `manual`, or `none`.

- `dhcp` - use NTP servers from DHCP
- `default` - use default NTP servers
- `manual` - use provided NTP servers, in this case at least one `<ntp-server>` entry must be specified
- `none` - NTP is disabled

If `source` is `dhcp`, `default`, or `none`, do not specify `<ntp-server>`.

<ntp-server> Required? No

Description: Specifies one or more NTP servers. To be used only with the `ntp` element and the attribute `manual`.

Attributes: None

<keymap> Required? No

Description: The name of the key map to use during installation. `<keymap>us</keymap>` The default value, `us` is considered if you do not specify a value for this element.

Attributes: None

<root-password> Required: No

Description: The desired root password for the XenServer host. If a password is not provided, a prompt is displayed when the host is first booted.

Attributes: You can specify a `type` that is either `hash` or `plaintext`

For example:

```
1 <root-password type="hash">hashedpassword</root-password>
```

The hashed value can use any hash type supported by `crypt(3)` in `glibc`. The default hash type is SHA-512.

You can use the following Python code to generate a hashed password string to include in the answer file:

```
1 python -c 'import crypt; print(crypt.crypt("mypasswordhere", crypt.mksalt(crypt.METHOD_SHA512)))'
```

<source> **Required:** Yes

Description: The location of the uploaded XenServer installation media or a Supplemental Pack. This element can occur multiple times.

Attributes: The attribute `type` can have one of the following values: `url`, `nfs`, or `local`.

If the value is `local`, leave the element empty. For example,

```
1 <source type="url">http://server/packages</source>
2 <source type="local" />
3 <source type="nfs">server:/packages</source>
```

<script> **Required:** No

Description: Where the post-install-script is located.

Attributes:

The attribute `stage` can have one of the following values: `filesystem-populated`, `installation-start`, or `installation-complete`.

- When the value `filesystem-populated` is used, the script runs just before the root file system is unmounted (for example, after installation/upgrade, initrds already built, and so on). The script receives an argument that is the mount point of the root file system.
- When the value `installation-start` is used, the script runs before starting the main installation sequence, but after the installer has initialized, and processed the answerfile. The script does not receive any arguments.
- When the value `installation-complete` is used, the script runs after the installer has finished all operations (and hence the root file system is unmounted). The script receives an argument that has a value of zero if the installation completed successfully, and is non-zero if the installation failed for any reason.

The attribute `type` can have one of the following values: `url`, `nfs`, or `local`.

If the value is `url` or `nfs`, put the URL or NFS path in the PCDATA. If the value is `local`, leave the PCDATA empty. For example,

```
1 <script stage="filesystem-populated" type="url">
2   http://prehost.example.com/post-install-script
3 </script>
4 <script stage="installation-start" type="local">
```

```
5     file:///scripts/run.sh
6 </script>
7 <script stage="installation-complete" type="nfs">
8     server:/scripts/installation-pass-fail-script
9 </script>
```

Note:

If a local file is used, ensure that the path is absolute. This generally means that the `file://` prefix is followed by another forward slash, and then the complete path to the script.

<admin-interface> **Required:** Sometimes

Note:

Required during install/reinstall but not during upgrade or restore.

Description: The single network interface to be used as the host administration interface.

Attributes:

Specify one of the following attributes:

- **name** - The name of your network interface, for example `eno12429np0`.
- **hwaddr** - The MAC address of your network interface, for example `00:00:11:aa:bb:cc`.

The attribute **proto** can have one of the following values: `dhcp` or **static**.

If you specify `proto="static"`, you must also specify all of these child elements:

Child elements

- **<ipaddr>**: The IP address
- **<subnet>**: The subnet mask
- **<gateway>**: The gateway

<timezone> **Required:** No

Description: The timezone in the format used by the TZ variable, for example Europe/London, or America/Los_Angeles. The default value is `Etc/UTC`.

<name-server> **Required:** No

Description: The IP address of a nameserver. Use one of these elements for each nameserver you want to use.

<hostname> **Required:** No

Description: Specify this element if you want to manually set a host name.

<ntp-server> **Required:** No

Description: Specify one or more NTP servers.

Start the network installation

After setting up the network servers required for a PXE boot installation, complete the following steps on the server that you are installing onto:

1. Start the system and enter the boot menu (**F12** in most BIOS programs).
2. Select to boot from your Ethernet card.
3. The system then PXE boots from the installation source you set up and the installation script starts.
 - If you have set up an answer file, the installation proceeds unattended.
 - If you have decided to do a manual installation, provide information when prompted. For more information, see [Install](#).

Boot from SAN

Boot-from-SAN environments offer several advantages, including high performance, redundancy, and space consolidation. In these environments, the boot disk is on a remote SAN and not on the local host.

The following types of boot-from-SAN configuration are supported:

- HBA and hardware Fibre Channel
- Software boot from iSCSI

Before setting up boot from SAN in XenServer, ensure that you have configured your hardware BIOS to enable this feature. Refer to your hardware vendor documentation for more information.

For a fully redundant boot-from-SAN environment, you must configure multiple paths for I/O access. For more information, see [Enable multipathing](#).

Note:

When using a boot-from-SAN environment it is highly recommended to enable sending logs to a remote syslog server. see [Send syslog messages to a central server](#). If remote syslog is not enabled and a system fault results in loss of access to the boot disk it can be very difficult to

determine the underlying cause and Technical Support may not be able to provide assistance without these logs.

HBA and hardware Fibre Channel

This type of boot-from-SAN deployment depends on SAN-based disk arrays with either hardware Fibre Channel or HBA iSCSI adapter support on the host. The host communicates with the SAN through a host bus adapter (HBA). The HBA's BIOS contains the instructions that enable the host to find the boot disk.

All of the configuration to set up boot from SAN through hardware Fibre Channel or an HBA adapter is done in your network infrastructure before you install XenServer on your servers. For information about how to complete this set up, see the documentation provided by the vendor.

Software boot from iSCSI

The software-boot-from-iSCSI feature enables customers to install and boot XenServer from SAN using iSCSI. Using this feature, XenServer can be installed to, booted from, and run from a LUN provided by an iSCSI target. The iSCSI target is specified in the iSCSI Boot Firmware Table. This capability allows the root disk to be attached through iSCSI. This boot disk can be located on the same target that provides an SR.

To use this feature, ensure that your environment meets the following requirements:

- The network interface or interfaces dedicated to iSCSI boot must be separate from the management interfaces and interfaces used for VM traffic.
- Storage (iSCSI targets) must be on a separate Layer 3 (IP) network to all other network interfaces with IP addresses on the host.
- Do not use tagged VLAN for the network interfaces dedicated to the iSCSI boot targets.
- We recommend that you have multipathing configured.

To configure the software-boot-from-iSCSI feature, you must add the `use_ibft` parameter to your boot parameters. How you add this parameter depends on your boot mode and the type of installation you are doing.

Enable the software-boot-from-iSCSI feature on a UEFI boot server during an installation from local media

1. Boot the computer from the installation media. For more information, see [Install the XenServer host](#).

Following the initial boot messages, you see a GRUB menu. This menu is shown for 5 seconds.

2. Use the cursor keys to select the **install** option:
3. Press the **e** key to edit the commands before booting.
4. Edit the line starting with the following:

```
1 module2 /EFI/xenserver/vmlinuz ...
```

Using the cursor keys, edit this line to include `use_ibft` at the end:

```
1 module2 /EFI/xenserver/vmlinuz ... use_ibft
```

5. Press **Enter**.
6. Continue your XenServer host installation process as normal.

Enable the software-boot-from-iSCSI feature on a UEFI boot server during a PXE boot installation When installing using PXE, ensure that you add the keyword **use_ibft** in the kernel parameters.

For example:

```
1 menuentry "XenServer Install (serial)" {
2
3     multiboot2 /EFI/xenserver/xen.gz dom0_max_vcpus=1-16 dom0_mem=max
      :8192M com1=115200,8n1 console=com1,vga
4     module2 /EFI/xenserver/vmlinuz console=hvc0 console=tty0
      answerfile_device=<interface_name> answerfile=http://<ip_address
      >/<path_to_answer_file> install use_ibft
5     module2 /EFI/xenserver/install.img
6 }
```

For more information about setting up PXE boot, see [Configure your TFTP server for PXE boot with UEFI](#).

Install supplemental packs

Supplemental packs are used to modify and extend the capabilities of XenServer by installing software into the control domain (dom0). For example, an OEM partner might want to ship XenServer with a set of management tools that require SNMP agents to be installed. You can install a supplemental pack either during initial XenServer installation, or any time afterwards on a running XenServer instance.

When installing supplemental packs during XenServer installation, unpack each supplemental pack into a separate directory on a web server.

You can install the supplemental pack in one of the following ways:

- During an interactive installation, when you are prompted to install supplemental packs, specify the URL to the supplemental pack media.

- If you are using an answer file for your install, add an additional `<source>` element to specify the location of the supplemental pack.

Upgrade from XenServer 8.4

April 27, 2026

By upgrading from an existing installation of XenServer 8.4 to XenServer 9, you can retain your existing VMs, SRs, and configuration.

Perform a rolling pool upgrade to keep all the services and resources offered by the pool available while upgrading all hosts in the pool. This upgrade method only takes one XenServer® host offline at a time. Critical VMs are kept running during the process by live migrating the VMs to other hosts in the pool.

You can complete a rolling pool upgrade in one of the following ways:

- If you have a Premium Edition license, you can use the XenCenter® **Rolling Pool Upgrade** wizard. This wizard organizes the upgrade path automatically and guides you through the upgrade procedure.

For more information, see [Rolling pool upgrade by using XenCenter](#).

- You can use the xe CLI to perform a rolling pool upgrade manually by live migrating running VMs between XenServer hosts accordingly.

For more information, see [Rolling pool upgrade by using the xe CLI](#).

Can I upgrade?

Ensure that you are able to upgrade to XenServer 9:

- Are your hosts currently running XenServer 8.4?

If not, you cannot upgrade directly to XenServer 9. Instead, perform a clean installation. For more information, see [Install](#).

If you are already using XenServer 9, do not attempt to update by using the installation ISO. Instead, apply updates to get your XenServer 9 pool to the latest level. For more information, see [Apply updates](#).

- Is your hardware compatible with XenServer 9?

Check that the hardware your pool is installed on is compatible with the version of XenServer you are about to upgrade to. For more information, see the [Hardware Compatibility List \(HCL\)](#).

- Are your VM operating systems supported by XenServer 9?

Check that the operating systems of your VMs are supported by XenServer 9. If your VM operating system is not supported, upgrade your VM operating system to a supported version before upgrading XenServer. For more information, see [Guest operating system support](#).

- Is your pool using GFS2 thin provisioned storage?

Check that the GFS2 cluster stack has been migrated to use version 3 of the Corosync cluster stack. See [GFS2 Cluster Version Management](#).

Before you start

Review the following information before starting your upgrade. Take the necessary steps to ensure that your upgrade process is successful.

Plan for the upgrade

1. Map your upgrade path carefully. Upgrading XenServer hosts, and particularly a pool of XenServer hosts, requires careful planning and attention to avoid losing any existing data.

Note the following information when planning your upgrade:

- You cannot migrate a VM from a newer version of XenServer to an older one.
 - Do not operate your pool in mixed-mode (with multiple versions of XenServer) for longer than is necessary. The pool operates in a degraded state during upgrade.
 - Key control operations are not available during the upgrade process. Do not attempt to perform any control operations.
 - Do not copy, shut down, or export VMs during the upgrade process.
 - Do not perform storage-related operations - such as adding, removing, or resizing virtual disks - during the upgrade process.
 - During the upgrade of the pool coordinator, the other hosts in the pool enter *emergency mode*.
2. Ensure that all hosts in your pool are synchronized to the correct date and time (for example, by using NTP).
 3. Ensure that your servers are not over-provisioned: check that servers have sufficient memory to carry out the upgrade.

Generally, if N equals the total number of servers in a pool, there must be sufficient memory across $N-1$ servers to run all live VMs in the pool. It is best to suspend any non-critical VMs during the upgrade process.

4. Ensure that your pool has shared storage to keep your VMs running during a rolling pool upgrade. If your pool does not have shared storage, you must suspend your VMs before upgrading because the VMs cannot be live migrated.

Storage live migration is not supported with rolling pool upgrades.

5. Note the following behaviors:

- The upgrade must use the same boot mode as the initial install.
- Boot-from-SAN settings are *not* inherited during the manual upgrade process. When upgrading using the ISO or PXE process, you must ensure that `multipathd` is correctly configured. For more information, see [Boot from SAN](#).
- When you upgrade XenServer, previously applied supplemental packs are removed and so they must be reapplied during or after the upgrade.

Prepare your pool

1. Take a backup of the state of your existing pool using the xe CLI command `xe pool-dump-database`.

Taking a backup of the state ensures that you can revert a partially complete rolling upgrade to its original state without losing VM data.

2. Disable high availability.

Prepare your VMs

1. If you have Windows VMs running in your pool, take the following steps for each VM:
 - Ensure that the latest version of the XenServer VM Tools for Windows is installed.
 - Take a snapshot of the VM.
2. If you have Linux VMs running in your pool, ensure that the latest version of the XenServer VM Tools for Linux is installed.
3. If you have NVIDIA vGPU-enabled VMs running on your pool, complete the following steps to migrate the pool while these VMs are running:
 - a) Ensure that the GPU you are using is supported on the version you plan to upgrade to.
 - b) Identify a version of the NVIDIA GRID drivers that is available for both your current version of XenServer and the version of XenServer you are upgrading to. If possible, choose the latest available drivers.
 - c) Install the new GRID drivers on your XenServer hosts and the matching guest drivers on any of your vGPU-enabled VMs.

- d) Ensure that you also have the version of the GRID driver that matches the XenServer version that you are upgrading to. You are prompted to install these drivers as a supplemental pack as part of the rolling pool upgrade process.
4. Empty the CD/DVD drives of all VMs in the pool.

Set up the installation media

1. If you are using XenCenter to upgrade your hosts, download and install the latest version of XenCenter from the [XenServer download site](#).

For more information, see [Install XenCenter](#).

2. Download the XenServer 9 installation ISO from the [XenServer download site](#). This installation ISO is regularly updated, ensure that you use the latest version rather than a previously downloaded ISO.

3. Prepare the installation media:

- To upgrade your hosts from a bootable USB, use a tool like [rufus](#) or [diskpart](#) to create a bootable USB by using the XenServer 9 installation ISO. Ensure that the tool does not alter the contents of the ISO file.
- To upgrade your hosts from a CD, burn the XenServer 9 installation ISO file to a CD.
- To upgrade your hosts from virtual media, go to the virtual console of your system and mount the XenServer installation ISO file as virtual media.
- To upgrade from a network location, configure your network location and installation files as described in [Host your installation media on NFS, FTP, or HTTP](#).

These steps are an overview of the process:

- a) If you want to boot the installer from TFTP, set up a network-accessible TFTP server and configure the installation files.
- b) Set up a network location where you can access the installation ISO through HTTP, FTP, or NFS.
- c) Unpack the installation ISO onto the network location. You can do this by mounting the ISO and copying its contents to the network location.
- d) Make note of the information you need during the upgrade:
 - For HTTP or FTP, make note of the URL for your HTTP or FTP repository, and a user name and password, if appropriate.
 - For NFS, make note of the server and path of your NFS share.

After these prerequisite steps are completed, you can perform a rolling pool upgrade by one of the following methods:

- Rolling pool upgrade by using XenCenter
- Rolling pool upgrade by using the xe CLI

Rolling pool upgrade by using XenCenter

The **Rolling Pool Upgrade** wizard guides you through the upgrade procedure and organizes the upgrade path automatically. For pools, each of the servers in the pool is upgraded in turn, starting with the pool

coordinator. Before starting an upgrade, the wizard conducts a series of pre-checks. These pre-checks ensure certain pool-wide features, such as high availability, are temporarily disabled and that each server in the pool

is prepared for upgrade. Only one server is offline at a time. Any running VMs are automatically migrated off each server before the upgrade is installed on that server.

Note:

The XenCenter **Rolling Pool Upgrade** wizard is only available if you have a Premium Edition license.

If you have not yet installed XenCenter, download the latest version from the [XenServer download site](#) and complete the steps in [Install XenCenter](#).

Note:

Use the latest XenCenter to upgrade from XenServer 8.4 to XenServer 9. Using an older version of XenCenter can result in a loss of connectivity.

Download the latest XenCenter from the [XenServer product downloads page](#).

The wizard can operate in manual or automatic mode:

- In manual mode, you must manually run the XenServer installer on each server in turn and follow the on-screen instructions on the serial console of the server. When the upgrade begins, XenCenter prompts you to insert the installation media or specify a network boot server for each server that you upgrade.
- In automatic mode, the wizard uses network installation files on an HTTP, NFS, or FTP server to upgrade each server in turn. This mode doesn't require you to insert installation media, manually reboot, or step through the installer on each server. If you perform a rolling pool upgrade in this manner, you must unpack the installation media onto your HTTP, NFS, or FTP server before starting the upgrade.

To upgrade XenServer hosts by using the XenCenter Rolling Pool Upgrade wizard:

1. On the XenCenter **Tools** menu, select **Rolling Pool Upgrade**.
2. Read the **Before You Start** information. Click **Next** to continue.
3. Select the pools and any individual hosts that you want to upgrade, and then click **Next**.
4. Choose one of the following modes:
 - **Automatic Mode** for an automated upgrade from network installation files on an HTTP, NFS, or FTP server. If using this method, ensure that you have completed all the necessary steps to set up the installation media for network install. For more information, see Set up the installation media.

If you choose **Automatic Mode** and are using IIS to host the installation media, ensure that double escaping is enabled on IIS before extracting the installation ISO on it.
 - **Manual Mode** for a manual upgrade either from a USB/CD/DVD or by using network boot (using existing infrastructure).

If you choose **Manual Mode**, you must run the XenServer installer on each host in turn. Follow the on-screen instructions on the serial console of the host. When the upgrade begins, XenCenter prompts you to insert the XenServer installation media or specify a network boot server for each host that you upgrade.
5. After you have selected your upgrade mode, click **Run Pre-checks**.
6. Follow the recommendations to resolve any upgrade pre-checks that have failed. If you want XenCenter to resolve all failed pre-checks automatically, click **Resolve All**.
7. When all pre-checks have been resolved, click **Next** to continue.
8. Prepare the XenServer installation media.
 - If you chose **Automatic Mode**, enter the installation media details. Choose **HTTP**, **NFS**, or **FTP** and then specify the URL, user name, and password, as appropriate.

Notes:

- If you choose FTP, ensure that you escape any leading slashes that are in the file path section of the URL.
 - Enter the user name and password associated with your HTTP or FTP server, if you have configured security credentials. Do not enter the user name and password associated with your XenServer pool.
 - XenServer supports FTP in passive mode only.
-
- If you chose **Manual Mode**, note the upgrade plan and follow the instructions.

9. Click **Start Upgrade**.
10. When the upgrade begins, the **Rolling Pool Upgrade** wizard guides you through any actions you must take to upgrade each host. Follow the instructions until you have upgraded and updated all hosts in the pools.
11. If you have vGPU-enabled VMs, when you reach the step that gives you the option to supply a supplemental pack, upload the NVIDIA driver that matches the one on your vGPU-enabled VMs. Ensure you upload the version of the driver for the XenServer version you are upgrading to.
12. The **Rolling Pool Upgrade** wizard prints a summary when the upgrade is complete. Click **Finish** to close the wizard.

Note:

If the upgrade or the update process fails for any reason, the **Rolling Pool Upgrade** wizard halts the process. This allows you to fix the issue and resume the upgrade or update process by clicking the **Retry** button.

After the upgrade

After your pool is upgraded, we recommend that you complete the following tasks:

- Configure updates and apply the latest set. For more information, see [Update your XenServer hosts](#).

After a rolling pool upgrade is complete, a VM might not be located on its home host. To relocate the VM, you can do one of the following actions:

- Live migrate the VM to its home host
- Shut down the VM and then start it on its home host

Rolling pool upgrade by using the xe CLI

Before performing a rolling pool upgrade by using the xe CLI, ensure that you have completed all the prerequisite steps in [Before you start](#).

Important:

Ensure that you upgrade all servers in your pool. We strongly advise against running a mixed-mode pool (one with multiple XenServer versions) for longer than necessary, as the pool operates in a degraded state during upgrade.

Key control operations are not available during the upgrade process. Do not attempt to perform any control operations. Though VMs continue to function as normal, VM actions other than mi-

grate are not available (for example, shut down, copy and export). In particular, it is not safe to perform storage-related operations such as adding, removing, or resizing virtual disks.

To perform a rolling pool upgrade by using the xe CLI:

Start with the pool coordinator:

1. Disable the pool coordinator. This prevents any new VMs from starting on or being migrated to the specified host.

```
1 xe host-disable host=<host_uuid>
```

Instead of the `host` parameter this command, you can use a host selector. For more information, see [Host selectors](#).

2. Ensure that no VMs are running on the pool coordinator. Shut down, suspend, or migrate VMs to other hosts in the pool.

- To shut down a VM, use the following command:

```
1 xe vm-shutdown
```

- To suspend a VM, use the following command:

```
1 xe vm-suspend
```

- To migrate a specific VM, use the following command:

```
1 xe vm-migrate
```

Migrating specified VMs to specified hosts gives you full control over the distribution of migrated VMs to other hosts in the pool.

- To evacuate the host, use the following command:

```
1 xe host-evacuate
```

Evacuating all VMs from a host leaves the distribution of migrated VMs to XenServer.

3. Shut down the pool coordinator by using the following xe CLI command:

```
1 xe host-shutdown
```

Important:

Always use the xe CLI or XenCenter to shut down or restart your hosts. Do not use other methods (such as hardware-based resets or operating system commands) to shut down or restart your hosts during normal operation. For more information, see [Shut down a XenServer host](#).

You are unable to contact the pool coordinator until the upgrade of the pool coordinator is complete. Shutting down the pool coordinator causes the other hosts in the pool to enter *emergency mode*. Hosts can enter emergency mode when they are in a pool whose pool coordinator has disappeared from the network and cannot be contacted after several attempts. VMs continue to run on hosts in emergency mode, but control operations are not available.

4. Boot the pool coordinator using the XenServer installation media and method of your choice (such as, USB or network).
5. Follow the XenServer installation procedure until the installer offers you the option to upgrade. Choose to upgrade.

When your pool coordinator restarts, the other hosts in the pool leave emergency mode and normal service is restored after a few minutes.

6. Start or resume any shutdown or suspended VMs.
7. Migrate any VMs that you want back to the pool coordinator.

If anything interrupts the upgrade of the pool coordinator or if the upgrade fails for any reason, do not attempt to proceed with the upgrade. Reboot the pool coordinator and restore to a working version.

Repeat these steps for all the other hosts in the pool:

1. Select the next XenServer host in your upgrade path. Disable the host.

```
1 xe host-disable host=<host_uuid>
```

Instead of the `host` parameter in this command, you can use a host selector. For more information, see [Host selectors](#).

2. Ensure that no VMs are running on the host. Shut down, suspend, or migrate VMs to other hosts in the pool.

- To shut down a VM, use the following command:

```
1 xe vm-shutdown
```

- To suspend a VM, use the following command:

```
1 xe vm-suspend
```

- To migrate a specific VM, use the following command:

```
1 xe vm-migrate
```

Migrating specified VMs to specified hosts gives you full control over the distribution of migrated VMs to other hosts in the pool.

- To evacuate the host, use the following command:

```
1  xe host-evacuate
```

Evacuating all VMs from a host leaves the distribution of migrated VMs to XenServer.

3. Shut down the host by using the following xe CLI command:

```
1  xe host-shutdown
```

Important:

Always use the xe CLI or XenCenter to shut down or restart your hosts. Do not use other methods (such as hardware-based resets or operating system commands) to shut down or restart your hosts during normal operation. For more information, see [Shut down a XenServer host](#).

4. Boot the host using the XenServer installation media and method of your choice (such as, USB or network).
5. Follow the XenServer installation procedure until the installer offers you the option to upgrade. Choose to upgrade.
6. After the host upgrade is complete, start or resume any shutdown or suspended VMs.
7. Migrate any VMs that you want back to the host.

If the upgrade of a subordinate host fails or is interrupted, you do not have to revert. Run the command `xe host-forget` in the pool to forget that host. Reinstall XenServer on the host, and then join it, as a new host, to the pool using the command `xe pool-join`.

After the upgrade

After your pool is upgraded, we recommend that you complete the following tasks:

- Configure updates and apply the latest set. For more information, see [Update your XenServer hosts](#).

After a rolling pool upgrade is complete, a VM might not be located on its home host. To relocate the VM, you can do one of the following actions:

- Live migrate the VM to its home host
- Shut down the VM and then start it on its home host

Other scenarios

Citrix Virtual Apps and Desktops environments

If you are using XenServer to host your Citrix Virtual Apps and Desktops workloads, see [Upgrade scenarios for Citrix Virtual Apps and Desktops](#).

Update your XenServer® hosts

June 3, 2026

XenServer 9 uses a faster, more efficient release process designed to deliver updates faster than ever. New features, bug fixes, and performance improvements are frequently pushed to the content delivery network (CDN) as available updates for your XenServer hosts and pools. Each set of updates is cumulative, meaning they include the latest features and fixes along with all prior updates, ensuring your environment remains optimized with minimal effort.

[This is an embedded video. Click the link to watch the video](#)

Update methods

Updates can be applied to your XenServer hosts and pools through the following methods:

- **Updates from the CDN:** Requires internet access to automatically download and apply updates from the CDN.
- **Offline updates:** Designed for air-gapped environments, allowing you to manually download and apply updates without internet access.

For advanced automation, you can also use the XenServer PowerShell module to script the upload and application of update bundles. For more information, see [Update your XenServer hosts by using the XenServer PowerShell module](#).

Updates from the CDN

Updates from the CDN provide the easiest way to keep your hosts and pools up-to-date. This method relies on an internet connection to download updates directly from the content delivery network (CDN). To ensure that you are always on the latest and greatest update, there is no picking and choosing - when you apply updates to your pool, it is updated to the latest, fully tested state.

How updates from the CDN work:

1. Updates are delivered through the secure content delivery network (CDN).
2. Configure your pool to automatically synchronize with the Early Access or Normal update channel. For more information about the update channels, see [Update channels and the CDN](#).
3. All available updates are downloaded to the pool coordinator during synchronization.
4. Use XenCenter® or the xe CLI to view, download, and apply these updates seamlessly.

For more information on configuring updates from the CDN using XenCenter, see [Updates from the CDN](#).

Updates from the CDN are suitable for environments with internet connectivity. For environments without internet access, use the Offline updates method.

Offline updates

File based licensing system EOL notification License Activation Service (LAS) is the next generation technology for product activations across the suite of Citrix® products.

Note:

File based licensing, traditionally used for activating various on-premises components, reached End of Life (EOL) on April 15, 2026.

As of April 15, 2026, LAS is the only way to activate and license Citrix on-premises components. To remain supported, your XenServer deployment must be on a LAS compatible version.

Requirement Should you have questions or concerns, contact Customer Care. Citrix may limit or suspend your Citrix Maintenance for non-compliance with these requirements without liability in addition to any other remedies Citrix may have at law or equity. These requirements don't apply where prohibited by law or regulation.

Offline updates are an alternative method for updating your XenServer hosts when internet connectivity is unavailable or restricted. Offline updates are relevant in the following scenarios:

- You want to update your hosts but you have XenCenter installed on an air-gapped machine.
If you have XenCenter installed on an air-gapped machine, you can still apply updates by leveraging a separate, internet-connected machine. Use this dedicated machine to download update bundles and then transfer the update bundles to your air-gapped XenCenter. You can upload and apply the update bundle from XenCenter on your isolated system, ensuring your hosts are up-to-date without compromising security.
- You want to join a new host to an existing pool, but your pool and host are on different update levels.

Typically, when adding a host to an existing pool, both the pool and the joining host must be updated to the latest level of updates. However, offline updates allow you to download and apply an update bundle to the joining host that matches the pool's current update state, simplifying the process of adding a host to a pool.

- You want to test updates before production rollout.

Offline updates offer you the flexibility to test specific updates in a controlled environment before rolling them out to production pools. This approach minimizes potential disruptions by allowing you to evaluate the impact of updates in a test pool first.

How offline updates work:

1. Download the latest update bundle from the following page using your dedicated internet-connected machine: [Download XenServer update bundles](#).

Each update bundle contains all previous updates. You can get to the required level by applying only one bundle.

2. Transfer the update bundle to a system that can communicate with your pool through XenCenter, the xe CLI, or the management API.
3. Use XenCenter, the xe CLI, or scripts to upload and apply the update bundle to keep your hosts up-to-date.

For more information on how to configure offline updates using XenCenter, see [Offline updates](#).

Update to pool level

Updating a standalone host to the same update level as an existing pool gives you an easier way to join a new host to that pool. This method is relevant in the following scenario:

- You want to join a new host to an existing pool, but your pool and host are on different update levels.

When adding a host to an existing pool, both the pool and the joining host must be updated to the same level of updates. This method allows you to synchronize the joining host with the pool and to apply updates stored on the pool coordinator that matches the pool's current update state, simplifying the process of adding a host to a pool.

How pool-level updates work:

1. Ensure the standalone host has the server certificate of the pool coordinator of the existing pool.
2. Configure the standalone host to use the existing pool as the update source.
3. Synchronize the standalone host with the pool to see the updates available to bring the host to the same level as the pool.

4. Use XenCenter or the xe CLI to apply the updates.

Update channels and the CDN

Regardless of the update method you choose, you must configure your XenServer host or pool with the appropriate update channel. The update channel determines how and when updates are made available to your hosts.

Updates are managed through two online channels and one offline option:

- **Early Access:** Updates are first pushed to the Early Access channel. This channel is ideal for test environments, allowing you to trial the latest updates as soon as they are released.

Note:

While Early Access updates are supported for production use, we do not recommend using this channel in critical production environments.

- **Normal:** Unless delayed, updates from the Early Access channel flow sequentially into the Normal channel on a regular schedule. This channel is recommended for production environments.
- **Offline:** Offline update bundles are only generated when updates are pushed to the Normal channel. When an update becomes available in the Normal channel, a corresponding update bundle is generated for Offline updates. These bundles can be downloaded from the following page and manually applied in air-gapped environments: [Download XenServer update bundles](#). Offline updates are not delivered through the CDN and do not require internet connectivity.

Note:

Occasionally, critical security patches or urgent fixes are made available simultaneously across all channels.

Lifecycle

During its lifecycle, XenServer 9 provides frequent updates that include new features, bug fixes, and security enhancements. These updates ensure your hosts remain secure and performant while enabling you to adopt improvements at the earliest opportunity.

You must periodically apply all available updates to maintain support and keep your systems aligned with the latest advancements. As a result, the behavior and feature set in XenServer 9 can change. For more information, see [XenServer lifecycle](#).

Support

In XenServer 9, new features, fixes, and improvements are frequently made available to your hosts. To remain in support, apply all outstanding updates within six months of their release, regardless of whether you use Automated Update or Offline Update methods. This regular update process ensures compatibility, performance, and security. If the update level of your pool is older than six months, we might ask you to reproduce the issue on the latest update level as part of our investigation.

Note:

Applying updates is only assured to work when the previous state of the pool is less than six months old. If you haven't applied updates for more than six months, updating to the latest update level is not tested and is not assured.

Get started with updates

For information on how to configure and apply updates for your XenServer hosts by using XenCenter, see [Apply updates by using XenCenter](#). Alternatively, you can use the xe CLI to apply updates to your XenServer hosts. For more information, see [Apply updates by using the xe CLI](#).

Important:

We do not support the direct usage or modification of the underlying update components in dom0. You can only use XenCenter or the xe CLI to configure and apply updates.

Apply updates by using XenCenter®

February 18, 2026

Keep your XenServer 9 hosts and pools up to date with the latest version of XenCenter, available from the [XenServer product downloads page](#).

Complete the following steps to update your XenServer® pools:

1. [Install the latest version of XenCenter](#).
2. [Install or upgrade to XenServer 9](#).
3. Configure updates for your pool.

You can configure updates for a pool or standalone host from the following locations:

- An update channel on the content delivery network (CDN). This method is recommended for internet-connected pools.

- An offline update from a downloaded bundle This alternative method is intended for use when internet connectivity to your XenServer hosts is unavailable or restricted.
 - An existing pool. This method enables you to bring a standalone host or a pool up to the same level of updates as another pool that is at a higher level of updates. It is recommended for joining a newly installed host to a pool that is not at the latest level of updates.
4. View available updates for your pool.
 5. Understand the guidance categories and update tasks.
 6. Apply updates to your pool.
 7. Complete any pending update tasks.

Configure updates for your pool

Updates from the CDN

Prerequisites

- Ensure that your XenServer® hosts have internet access and can connect to the required update domains.
- If your hosts are behind a firewall, allow access to subdomains of ops.xenserver.com. For more information, see [Connectivity requirements](#).

Note:


Applying updates is only assured to work when the previous state of the pool is less than six months old. If you haven't applied updates for more than six months, updating to the latest update level is not tested and is not assured.

Follow these steps to configure your hosts for applying updates from the CDN:

1. In XenCenter®, on the **Tools** menu, select **Configure Updates**. Alternatively, go to the **Updates** section under your pool's **General** tab and select **Configure Updates** or right-click on your pool and select **Updates > Configure Updates**. The **Configure Server Updates** window opens.
2. On the **XenServer 8.4** tab, select the pools or hosts that you want to configure.
3. Under **Update Channel**, specify how soon you want to access updates. Your pool or host can be subscribed to one of the following update channels:
 - **Early Access**
 - **Normal**
4. Under **Synchronization Schedule**, select how often you want your XenServer pool to synchronize with the update channel. This can be daily, or weekly on a certain day of the week.

Note:

- After synchronizing, apply the updates to your pool as soon as possible to benefit from the latest updates.
 - If you designate a new pool coordinator after synchronizing but before applying updates to the hosts in the pool, you must synchronize again with the new pool coordinator before you can update the pool.
 - Do not synchronize your XenServer pool while the pool is in the process of being updated.
5. (Optional) Under **Proxy Server**, specify a proxy to use when synchronizing with the update channel. This proxy server is used for communication between the host and the public CDN.
 6. Click **Apply** to apply the configuration changes to your XenServer pool and then repeat the above steps to configure updates for the rest of your XenServer pools.
 7. If you are happy with the configuration changes to your pools, click **OK** to save your changes and close the **Configure Server Updates** window. When you first set up your host or pool with an update channel (or if you later change your host or pool to synchronize with a different update channel), you are prompted to synchronize your host or pool with the update channel immediately. In the dialog box that opens, select **Yes, Synchronize** if you want to synchronize your host or pool with the update channel immediately.



Alternatively, you can choose to synchronize your host or pool with the update channel at a later date. To synchronize your host or pool later, go to the **Updates** section under your pool's **General** tab and select **Synchronize Now** or right-click on your pool and select **Updates > Synchronize with update channel**.

8. As soon as your pool synchronizes with the update channel, apply the downloaded updates to your pool by using the **Install Updates** wizard. For more information, see [Apply updates to your pool](#).

After configuring your XenServer pool or host, you can find the following information in XenCenter, in the **Updates** section under the **General** tab of your pool or host:

- The update channel you are subscribed to
- The last time your pool synchronized with the update channel
- The checksum of that synchronization
- The checksum of the last update you applied

Offline Updates

Prerequisites

- Your host must be installed with the ISO published on October 7, 2024, or later.
- You must use XenCenter version 2024.4.0 or later to apply update bundles.

Follow these steps to configure Offline Updates:

1. Download the latest update bundle from the following page using your dedicated internet-connected machine: [Download XenServer update bundles](#).
2. Transfer the update bundle to a location accessible from XenCenter.
3. In XenCenter, on the **Tools** menu, select **Configure Updates**. Alternatively, go to the **Updates** section under your pool's **General** tab and select **Configure Updates** or right-click on your pool and select **Updates > Configure Updates**. The **Configure Server Updates** window opens.
4. On the **XenServer 8.4** tab, select the pools or hosts that you want to configure.
5. Under **Update Channel**, select **Offline** to configure your host or pool to use the **Offline** update channel. Click **OK**.
6. You are prompted to upload an update bundle to your pool immediately.

To upload an update bundle immediately, click **Yes** and in the window that opens, navigate to the location of the downloaded update bundle `.xsbundle` file on your computer and select it. Click **Next**.

Alternatively, you can choose to upload an update bundle at a later date. To upload an update bundle later, go to the **Updates** section under your pool's **General** tab and select **Upload bundle** or right-click on your pool and select **Updates > Upload bundle**. The **Select update bundle file** window opens and you can select your `.xsbundle` file. You can also upload an update bundle using the **Install Updates** wizard. For more information, see [Apply updates to your pool](#).

7. Review available updates for your pool and check if there are any update tasks required before or after applying updates. For more information, see [View available updates for your pool](#).
8. Apply the updates to your pool by using the **Install Updates** wizard. For more information, see [Apply updates to your pool](#).

After configuring your XenServer pool or host, you can find the following information in XenCenter, in the **Updates** section under the **General** tab of your pool or host:

- The update channel your XenServer pool or host is configured to use
- The last time you uploaded an update bundle
- The checksum of that update bundle
- The checksum of the last update you applied

Update from a remote pool

This method enables you to update a pool or standalone host to the same level of updates as another pool (referred to as the “remote pool”). A remote-pool update synchronizes the host or pool with the remote pool and updates from update files held on the remote pool coordinator. Use this process to more easily join a new host to an existing pool by matching that pool’s update level.


For a remote pool to be used as an update channel by a pool or standalone host, it must meet the following prerequisites:

- The remote pool must be at a more recent level of updates than the standalone host.
- The remote pool must have the latest synchronized updates applied to it.
- The pool coordinator has not changed since the remote pool was last updated.

Configure a remote pool as an update channel

Follow these steps to configure your host or pool to use another pool as an updates channel:

1. In XenCenter, on the **Tools** menu, select **Configure Updates**. Alternatively, go to the **Updates** section under your pool’s **General** tab and select **Configure Updates** or right-click on your pool and select **Updates > Configure Updates**. The **Configure Server Updates** window opens.
2. On the **XenServer 8.4** tab, select the pools or hosts that you want to configure.
3. Under **Update Channel**, select the remote pool you want to synchronize with from the section **Channels on other pools or standalone servers**. Click **OK**.
4. Click **Apply** to apply the configuration changes to your XenServer pool and then repeat the above steps to configure updates for the rest of your XenServer pools.
5. If you are happy with the configuration changes to your pools, click **OK** to save your changes and close the **Configure Server Updates** window. When you first set up your host or pool with an update channel (or if you later change your host or pool to synchronize with a different update channel), you are prompted to synchronize your host or pool with the update channel immediately.
6. In the dialog box that opens, select **Yes, Synchronize** if you want to synchronize your host or pool with the update channel immediately.



Alternatively, you can choose to synchronize your host or pool with the update channel at a later date. To synchronize your host or pool later, go to the **Updates** section under your pool's **General** tab and select **Synchronize Now** or right-click on your pool and select **Updates > Synchronize with update channel**.

7. As soon as your pool synchronizes with the update channel, apply the downloaded updates to your pool by using the **Install Updates** wizard. For more information, see [Apply updates to your pool](#).

After configuring your XenServer pool or host, you can find the following information in XenCenter, in the **Updates** section under the **General** tab of your pool or host:

- The update channel you are subscribed to
- The last time your pool synchronized with the update channel
- The checksum of that synchronization
- The checksum of the last update you applied

View available updates for your pool

XenCenter issues notifications about available updates for your hosts and pools under the **Updates** tabs in the **Notifications** view. The **Updates** tabs are split into XenServer 8.4 updates and Citrix Hy-

pervisor updates.

The XenServer 8.4 **Updates** tab refreshes in the following scenarios:

- For hosts and pools subscribed to the **Early Access** or **Normal** update channel, the tab refreshes based on the synchronization schedule you've set (either daily or weekly on a specific day).

Alternatively, you can fetch the latest available updates for your hosts and pools by manually synchronizing with the **Early Access** or **Normal** update channels.

- For hosts and pools configured for Offline Updates, the tab refreshes each time you upload an update bundle to the host or pool.

Alternatively, you can fetch the latest available updates for your hosts and pools by uploading the latest update bundle to your host or pool.

To view the latest updates available for pools configured to apply updates from the CDN, synchronize your XenServer pool with the update channel. You can do this in XenCenter® from the following places:

- In the **By Server** view of the **Updates** tab, you can choose to **Synchronize All** to synchronize all pools managed by XenCenter or **Synchronize Selected** to synchronize the selected pools.
- Alternatively, go to the **Updates** section under your pool's **General** tab and select **Synchronize Now** or right-click on your pool and select **Updates > Synchronize with update channel**.

To view the latest available updates for pools configured for Offline Updates, upload the latest update bundle to your pool. You can do this in XenCenter from the following places:

- In the **By Server** view of the **Updates** tab, you can choose to **Upload Bundle to All** to upload a bundle to all pools managed by XenCenter that use the **Offline** channel or **Upload Bundle to Selected** to upload a bundle to the selected pools.
- Alternatively, go to the **Updates** section under your pool's **General** tab and select **Upload bundle** or right-click on your pool and select **Updates > Upload bundle**.

You can then review all updates available for your XenServer 8.4 pools. The types of updates are:

- **Security fixes**
- **Bug fixes**
- **Improvements**
- **New features**
- **Preview features**
- **Foundational changes**

Note:

Foundational changes are non-customer visible groundwork changes to maintain and improve the product.

In the main panel, use the **View** option to choose whether the updates are displayed **By Server** or **By Update**.

By server

Updates are grouped by host and type of update.

Updates (27)Filters are OFF

View | Filter by: Server | Configure Updates... | Synchronize All | Install Updates... | Export

Updates available at last synchronization	Update channel	Last synchronized	Last updated
Pool-A xrtmia-05-22 No updates found at last synchronization xrtmia-05-32	Early Access	Mar 5, 2024 8:58:49...	Never
Pool-B lcy2-dt61 The server will be evacuated prior to installing updates Mandatory update tasks: Reboot the server 2 security fixes 11 bug fixes 5 improvements 3 new features New Feature: Add multipath check to NRPE integration New Feature: Monitor host and dom0 resources with SNMP New Feature: Supporting Linux guests UEFI and Secure Boot 6 foundational changes 73 packages will be updated	Early Access	Mar 5, 2024 9:26:08...	Never

New Feature: Supporting Linux guests UEFI and Secure Boot

Supporting Linux guests UEFI and Secure boot.
Current UEFI-supported Linux distributions include:
* Red Hat Enterprise Linux 8
* Red Hat Enterprise Linux 9 (preview)
* Ubuntu Focal Fossa 20.04
* Ubuntu Jammy Jellyfish 22.04

New feature

<https://docs.xenserver.com/en-us/xenserver/8/vms>

XENSER-2024-980

You can filter the update information by server. Select an update and hover over it to view detailed information about the update.

To view this information about your available updates offline, select **Export All** to export the information as a `.md` file. The contents of the `.md` file are grouped by pool and then host. For each host, the

file lists the following information:

- Any update tasks for this host or its VMs. For more information about mandatory, recommended, and full effectiveness tasks, see [Understand the guidance categories and update tasks](#).
- Whether these updates contain live patches that can be applied to this server
- Updates grouped by update type
 - The update name
 - A description of the update
- A list of the RPMs to be updated on the host

By update

All updates are listed chronologically in order of release.

The screenshot shows the 'Updates (27)' window in XenServer. The window has a header bar with 'Updates (27)' and 'Filters are OFF'. Below the header is a toolbar with 'View', 'Filter by: Server', 'Update Type', 'Update Task', 'Install Updates...', and 'Export'. A 'Columns' dropdown is on the right. The main area is a table with columns: 'Update', 'Server / Pool', 'Mandatory tasks', 'Recommended...', and 'Full effectiveness...'. The table lists various updates, including 'Updates to storage manager', 'Toolstack fixes and improvements', 'Security Fix: Security update to the kernel', 'Update fastling driver to version 8.74.0.', 'Updates to Xen', 'Update to VNCterm', 'New Feature: Add multipath check to M integration', 'GFS2 SR Improvements', 'Update to openssh', 'Updates to Storage Manager', 'Toolstack fixes', 'Update to Xen 4.17.3', 'Update cgroup libraries', 'New Feature: Monitor host and dom0 resources with SNMP', 'GFS2 Storage improvements', 'Fix SR-IOV on Intel E810 NICs', 'Packaging changes for pvsproxy', 'Security Fix: Security Update for Xen', 'Updates to xenserver-release', 'New Feature: Supporting Linux guests UEFI and Secure Boot', 'Build updates for qemu', and 'Build changes for xenserver-release'. A tooltip is displayed over the 'Security Fix: Security update to the kernel' row, showing details: 'Security Fix: Security update to the kernel', '* Fix for XSA-448 / CVE-2023-46838', 'Security fix', 'Severity: High', 'Mandatory update tasks: Reboot the server', and 'XENSER-2024-1009'.

Update	Server / Pool	Mandatory tasks	Recommended...	Full effectiveness...
Updates to storage manager	Pool-B			
Toolstack fixes and improvements	Pool-B			
Security Fix: Security update to the kernel	Pool-B			
Update fastling driver to version 8.74.0.	Pool-B			
Updates to Xen	Pool-B			
Toolstack fixes and improvements	Pool-B			
Update to VNCterm	Pool-B			
New Feature: Add multipath check to M integration	Pool-B			
GFS2 SR Improvements	Pool-B			
Update to openssh	Pool-B			
Updates to Storage Manager	Pool-B			
Toolstack fixes	Pool-B			
Update to Xen 4.17.3	Pool-B	Reboot Server		
Update cgroup libraries	Pool-B			
New Feature: Monitor host and dom0 resources with SNMP	Pool-B			
GFS2 Storage improvements	Pool-B			
Fix SR-IOV on Intel E810 NICs	Pool-B			
Packaging changes for pvsproxy	Pool-B			
Security Fix: Security Update for Xen	Pool-B	Reboot Server		
Updates to xenserver-release	Pool-B			
New Feature: Supporting Linux guests UEFI and Secure Boot	Pool-B			
Build updates for qemu	Pool-B			
Build changes for xenserver-release	Pool-B			

You can filter the update information by the server it can be applied to, by the update type, and by any update tasks that apply to it. Select an update and hover over it to view detailed information about the update.

To view this information about your available updates offline, select **Export All** to export the information as a `.csv` file. The `.csv` file contains the following information:

- The update type
- The update name
- The servers that this update can be applied to
- Whether the update includes a live patch that can be applied to the servers
- The mandatory, recommended, and full effectiveness tasks

For more information about mandatory, recommended, and full effectiveness tasks, see Understand the guidance categories and update tasks.

- To apply the updates to your hosts or pools, select **Install Updates** to open the **Install Updates** wizard. For more information, see the following section Apply updates to your pool.

Understand the guidance categories and update tasks

Some tasks (such as evacuating or rebooting your hosts) might be required before and after applying updates to your pool. Sometimes, there are no update tasks required.

Guidance categories

XenServer® tries to minimize the disruption to your VMs that these update tasks might cause by categorizing the tasks into **Mandatory**, **Recommended**, **Full-effectiveness**, and **Live patch**. These categorizations enable you to judge whether an update task that might cause downtime or minor disruption for your hosts or VMs is necessary for your environment and risk profile.

Updates can have tasks listed in more than one of these categories. For example, an update might require that you restart the host to get the full-effectiveness of the update, but recommend restarting the toolstack to get most of the benefit of the update with less potential disruption to the pool.

During the update process, you can choose to carry out one of the following three levels of tasks:

1. Mandatory
2. Mandatory + Recommended
3. Mandatory + Recommended + Full-effectiveness

Mandatory Mandatory tasks *must* be performed after an update otherwise the system might fail at runtime. These actions are required to enable critical fixes and ensure that your environment is secure and stable. When you apply updates, XenCenter® performs these tasks. You cannot opt out of mandatory tasks.

Recommended Recommended tasks are the tasks that we recommend you perform to get the benefit of the majority of features and fixes delivered in the updates. When you apply updates, these tasks are selected by default in XenCenter, but you can opt out of performing them. If you choose not to perform these tasks now, they are listed in the pending tasks for the applicable pool, host, or VM.

Why perform the recommended tasks:

- These tasks are the ones that ensure a secure, stable XenServer environment.

Why opt out of the recommended tasks:

- After reviewing the detailed information for the updates, you judge that the risk of not fully applying these updates now is acceptable.
- The recommended tasks cause unwanted disruption to your VMs now.

Full-effectiveness Full effectiveness tasks are required to get the benefit of the related update. The updates that have full-effectiveness tasks associated with them are usually relevant only to users on certain hardware or using specific features.

Review the update information to understand whether these tasks are required for your environment. When you apply updates, these tasks are not selected by default in XenCenter, but you can choose to perform them during the update if you believe the update applies to your environment or configuration. If you choose not to perform these tasks now, they are listed in the pending tasks for the applicable pool, host, or VM.

Why perform the full-effectiveness tasks:

- The updates that have full-effectiveness tasks are relevant to your hardware, environment, or configuration.

Why remain opted out of the full-effectiveness tasks:

- The updates that have full-effectiveness guidance are not relevant to your hardware, environment, or configuration.
- The full-effectiveness tasks cause unwanted disruption to your VMs now.
- You do not need the benefits of these updates right now.

If the full-effectiveness tasks apply to your environment, but you have opted to defer them, plan to complete these tasks during a suitable maintenance window to maintain the stability of your environment.

Live patches Updates to certain components can include a live patch. A live patch enables updates to be applied to core components without requiring a host reboot.

Whether a live patch can be applied to your hosts depends on the version of the component that was installed when the hosts were last rebooted. If an update can be applied as a live patch to your hosts, the live patch guidance replaces the recommended guidance.

Example:

You have two pools. Pool A is updated to a recent level. Pool B has not been updated for some time. We release a new update that has the recommended update task “Restart host” and the live patch update task “Restart toolstack”.

In pool A, the live patch can be applied to these more up-to-date hosts. XenCenter recommended guidance shows “Restart toolstack”. The less disruptive task from the live patch guidance overrides the recommended guidance.

In pool B, the live patch cannot be applied to the hosts as they are at an older level. XenCenter recommended guidance shows “Restart host”. The recommended guidance remains applicable. The live patch guidance is irrelevant in this case.

Sometimes only some of the fixes in an update are enabled when the update is applied as a live patch. In this situation, you might still see a recommended reboot for an update even though the live patch applies. Review the update details to understand whether you need all fixes in the update or just those fixes enabled by the live patch. You can then use this information to choose whether to perform the tasks. For more information, see [View available updates for your pool](#).

Update tasks

Applying updates might require you to complete one or more tasks, either before or after the updates are applied. These tasks can appear in any guidance category.

Update tasks for your host Tasks required *before* applying updates:

- **Evacuate server:** All VMs must be migrated off the XenServer host or shutdown before applying the update. To complete this task, XenCenter migrates any VMs off the host. While this task is in progress, the XenServer pool is operating at reduced capacity as one host is temporarily unavailable to run VMs. This task is sometimes carried out as part of the ‘Reboot host’ task.

Tasks required *after* applying updates:

- **Reboot server:** The XenServer host must be restarted. Any VMs are migrated off the host and the host is restarted. While this task is in progress, the XenServer pool is operating at reduced capacity as one host is temporarily unavailable to run VMs.

- **Restart toolstack:** The toolstack on the host must be restarted. If performed on the pool coordinator, the connection to the pool is temporarily lost. Restarting on other hosts has no visible effect on the pool's state.

Update tasks for your VM Some updates provide new features for your VMs. These updates might require the following tasks on your VMs:

- **Reboot VM:** The VM must be restarted. In XenCenter, the VM displays a red stop icon (square on red) while restarting. Once the task is complete, the green play icon appears, indicating the VM is running. During the reboot, the VM is unavailable to end users.
- **Restart device model:** The device model for VMs on the updated host must be restarted. In XenCenter, the VM shows a yellow play triangle while the device model restarts. Once complete, the green play icon appears. During this time, you cannot stop, start, or migrate the VM. The end user of the VM might experience a brief pause and resume in their session. To support this task on a Windows VM, the VM must have XenServer VM Tools for Windows installed.

Apply updates to your pool

The update installation mechanism in XenCenter® applies the updates to your hosts and pools using the **Install Updates** wizard. During the process, XenCenter automatically works out the least impactful action required after applying all available updates. The **Install Updates** wizard automatically performs these steps:

1. If required, it migrates VMs off each host.
2. If required, it places the host in Maintenance mode.
3. It applies the updates.
4. If required, it runs any necessary update tasks such as rebooting the host, restarting the toolstack, or rebooting the VMs.
5. It migrates the VMs back to the updated host.

Any actions that are taken at the pre-check stage to enable the updates to be applied, such as turning off high availability, are reverted.

Before you start

Before you apply an update to your servers, pay careful attention to the following:

- Back up all your servers
- If the set of updates includes updates to a driver, ensure that the new version of the driver is compatible with your hardware and firmware. For more information, see the [Hardware Compatibility List \(HCL\)](#) and your hardware vendor documentation.

- Ensure that High Availability (HA) is not enabled on any pool where you intend to apply updates.
- Ensure that you are logged in to XenCenter® as a Pool Administrator or Pool Operator, or using a local root account.

Install updates

Updates from the CDN

The following section provides step-by-step instructions on applying updates to your internet-connected pools using the **Install Updates** wizard:

1. From the XenCenter® menu, select **Tools** and then **Install Updates**.
2. In the **Install Updates** wizard, select **XenServer 8.4** and review the information on the **Before You Start** page. Click **Next** to continue.
3. Select **Apply updates**. Click **Next** to continue.
4. Select the XenServer® pools or hosts that you want to update. Click **Next** to progress to the next wizard page.
5. After updates are applied, some tasks, such as rebooting your hosts, might be required. On the **Update Mode** page, select the level of update tasks that you want XenCenter to automatically carry out after applying updates to your pool.

Tasks are categorized as **Mandatory**, **Recommended**, or **Full Effectiveness** under **Tasks**. If there are no update tasks required, the page displays a note saying **No action required**.

By default, XenCenter chooses the recommended level of update tasks. However, mandatory tasks are always performed and cannot be deselected.

To review these tasks in detail, minimize the wizard and go to the **Notifications > Updates** tab. For more information about the different types of update tasks and the guidance levels provided by XenCenter, see *Understand the guidance categories and update tasks*.

Click **Next** to progress to the next wizard page and begin pre-checks.

6. The wizard performs several pre-checks to verify that the updates can be applied on your host or pool. For example, you must have synchronized your host or pool with the update channel within the past week.

Follow the on-screen recommendations to resolve any pre-checks that have failed. If you prefer XenCenter to automatically resolve all failed pre-checks, select **Resolve All**. When the pre-checks have been resolved, select **Next** to continue.

7. Click **Install updates** for XenCenter to begin installing updates for your host or pool.

8. The **Install Updates** wizard shows the progress of the update, displaying the major operations that XenCenter performs while updating each host in the pool.

Notes:

- If the update process cannot complete for any reason, XenCenter halts the process. This halt allows you to fix the issue and resume the update process by clicking the **Retry** button.
- If you select **Cancel** at this stage, the **Install Updates** wizard reverts any changes that you might have made to resolve failed pre-checks (such as pausing High Availability and Workload Balancing or suspending VMs).

Click **Finish** to complete the updates and close the **Install Updates** wizard.

Offline Updates

The following section provides step-by-step instructions for applying updates to pools configured for offline updates using the **Install Updates** wizard:

1. From the XenCenter menu, select **Tools** and then **Install Updates**.
2. In the **Install Updates** wizard, select **XenServer 8.4** and review the information on the **Before You Start** page. Click **Next** to continue.
3. If you have previously uploaded an update bundle, select **Apply updates**. Otherwise, choose **Select an update bundle or supplemental pack from disk** and click **Browse** to navigate to the location of the downloaded update bundle . `xsbundle` file on your computer.

Note:

You can download the latest update bundle from the following page using your dedicated internet-connected machine: [Download XenServer update bundles](#). You then transfer the update bundle to a location accessible from XenCenter.

Click **Next** to continue.

4. Select the XenServer pools or hosts that you want to update. Click **Next**.
5. If you selected to upload an update bundle, the **Install Updates** wizard displays the progress of uploading the bundle to your pool or hosts on the **Upload** page. However, if you are applying a previously uploaded update bundle, this page is not displayed.

Note:

If you've just uploaded an update bundle that is older than the one currently applied to your pool, XenCenter displays that there are no available updates to apply and you cannot

proceed further.

Once the process is complete, click **Next** to continue.

6. After updates are applied, some tasks, such as rebooting your hosts, might be required. On the **Update Mode** page, select the level of update tasks that you want XenCenter to automatically carry out after applying updates to your pool.

Tasks are categorized as **Mandatory**, **Recommended**, or **Full Effectiveness** under **Tasks**. If there are no update tasks required, the page displays a note saying **No action required**.

By default, XenCenter chooses the recommended level of update tasks. However, mandatory tasks are always performed and cannot be deselected.

To review these tasks in detail, minimize the wizard and go to the **Notifications > Updates** tab. For more information about the different types of update tasks and the guidance levels provided by XenCenter, see [Understand the guidance categories and update tasks](#).

Click **Next** to progress to the next wizard page and begin pre-checks.

7. The wizard performs several pre-checks to verify that the updates can be applied on your host or pool.

Follow the on-screen recommendations to resolve any pre-checks that have failed. If you prefer XenCenter to automatically resolve all failed pre-checks, select **Resolve All**. When the pre-checks have been resolved, select **Next** to continue.

8. Click **Install updates** for XenCenter to begin installing the updates to your host or pool.
9. The **Install Updates** wizard shows the progress of applying the updates, displaying the major operations that XenCenter performs while updating each host in the pool.

Notes:

- If the update process cannot complete for any reason, XenCenter halts the process. This halt allows you to fix the issue and resume the update process by clicking the **Retry** button.
- If you select **Cancel** at this stage, the **Install Updates** wizard reverts any changes that you might have made to resolve failed pre-checks (such as pausing High Availability and Workload Balancing or suspending VMs).

Click **Finish** to complete the updates and close the **Install Updates** wizard.


Complete any pending update tasks

After applying updates, complete any pending update tasks to ensure your hosts and VMs operate as expected.

Note:

Before applying updates, you can review the required update tasks on the XenServer 8.4 **Updates** tab in the **Notifications** view. For more information, see View available updates for your pool.


If you skip any tasks during the update process, the pending tasks for each pool, host, or VM appear in the **Updates** section of the **General** tab for the pool, host, or VM.



Pool pending tasks

This section shows the pending tasks for all hosts in the pool.


It also shows a checksum that indicates the level that the current pool coordinator has synchronized to and checksums for each host that indicate the level of the updates installed. These checksums can provide useful information if you need to contact Technical Support.



Host pending tasks

This section shows the pending tasks for the XenServer® host.

It also shows a checksum that indicates the level of the updates installed. This checksum can provide useful information if you need to contact Technical Support.



VM pending tasks

This section shows the pending tasks for a VM.

Apply updates by using the xe CLI

April 2, 2026

Apply updates to your XenServer 9 hosts and pools by using the xe CLI.

Complete the following steps to update your XenServer® pools:

1. Configure updates for your pool.

You can configure updates for a pool or standalone host from the following locations:

- An update channel on the content delivery network (CDN). This method is recommended for internet-connected pools.
- An offline update from a downloaded bundle This alternative method is intended for use when internet connectivity to your XenServer hosts is unavailable or restricted

In addition, you can configure updates for a standalone host from the following location:

- An existing pool. This method enables you to bring a standalone host up to the same level of updates as a pool. It is recommended for joining a newly installed host to a pool that is not at the latest level of updates.
2. Synchronize your pool with your configured update channel.
 3. View available updates for your pool.
 4. Understand the guidance categories and update tasks.
 5. Apply updates to your pool.
 6. Complete any pending update tasks.

Updates from the CDN

To update from the CDN, your environment must meet the following prerequisites:

- Ensure that your XenServer hosts have internet access and can connect to the required update domains.
- If your hosts are behind a firewall, allow access to subdomains of ops.xenserver.com. For more information, see [Connectivity requirements](#).

Note:

Applying updates is only assured to work when the previous state of the pool is less than six

months old. If you haven't applied updates for more than six months, updating to the latest update level is not tested and is not assured.

Configure a CDN updates channel

Follow these steps to configure your hosts for applying updates from the CDN:

1. Create and enable the Early Access update channel for your pool:

```
1 pool_uuid=$(xe pool-list --minimal)
2
3 base_binary_url="https://repo.ops.xenserver.com/xs8/base"
4 base_source_url="https://repo-src.ops.xenserver.com/xs8/base"
5 base_repo_uuid=$(xe repository-introduce name=label=base_repo name
  -description=Base binary-url=<base_binary_url> source-url=<
    base_source_url> update=false)
6
7 update_binary_url="https://repo.ops.xenserver.com/xs8/earlyaccess"
8 update_source_url="https://repo-src.ops.xenserver.com/xs8/
  earlyaccess"
9 update_repo_uuid=$(xe repository-introduce name=label=
  early_access_repo name-description="Early Access" binary-url=<
    update_binary_url> source-url=<update_source_url> update=true)
10
11 xe pool-param-set uuid=$pool_uuid repositories=<base_repo_uuid>,<
  update_repo_uuid>
```

Alternatively, create and enable the Normal update channel for your pool:

```
1 pool_uuid=$(xe pool-list --minimal)
2
3 base_binary_url="https://repo.ops.xenserver.com/xs8/base"
4 base_source_url="https://repo-src.ops.xenserver.com/xs8/base"
5 base_repo_uuid=$(xe repository-introduce name=label=base_repo name
  -description=Base binary-url=<base_binary_url> source-url=<
    base_source_url> update=false)
6
7 update_binary_url="https://repo.ops.xenserver.com/xs8/normal"
8 update_source_url="https://repo-src.ops.xenserver.com/xs8/normal"
9 update_repo_uuid=$(xe repository-introduce name=label=normal name-
  description="Normal" binary-url=<update_binary_url> source-url
    =<update_source_url> update=true)
10
11 xe pool-param-set uuid=$pool_uuid repositories=<base_repo_uuid>,<
  update_repo_uuid>
```

2. Retrieve a list of the currently enabled repository UUIDs:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-param-get uuid=$pool_uuid param-name=repositories
```

3. Using the repository UUID, view more details about a particular repository:

```
1 xe repository-param-list uuid=<UUID>
```

4. (Optional) Configure and enable an HTTP connect proxy server that is used for communication between the host and the public CDN that hosts the repositories:

```
1 xe pool-configure-repository-proxy proxy-url=<http://proxy.example.com> proxy-username=<proxy-user> proxy-password=<proxy-password>
```

Disable the proxy server configuration:

```
1 xe pool-disable-repository-proxy
```

View the proxy server configuration:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-param-get uuid=$pool_uuid param-name=repository-proxy-url
3 xe pool-param-get uuid=$pool_uuid param-name=repository-proxy-username
```

Synchronize new updates for your pool

Enable your pool to automatically synchronize with the update channel by configuring a synchronization schedule. You can schedule a synchronization to take place daily or weekly on a certain day of the week. Synchronizing your pool with the update channel downloads all available updates to the pool coordinator and you can then apply all downloaded updates to your pool.

1. Set your pool to synchronize daily:

```
1 xe pool-configure-update-sync update-sync-frequency=daily update-sync-day=0
2 xe pool-set-update-sync-enabled value=true
```

Alternatively, set your pool to synchronize weekly:

```
1 xe pool-configure-update-sync update-sync-frequency=weekly update-sync-day=1 (# 0 is Sunday, 1 is Monday, etc)
2 xe pool-set-update-sync-enabled value=true
```

2. View your synchronization configuration:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-param-get uuid=$pool_uuid param-name=update-sync-frequency
3 xe pool-param-get uuid=$pool_uuid param-name=update-sync-day
4 xe pool-param-get uuid=$pool_uuid param-name=update-sync-enabled
```

3. Get the timestamp of your pool's last successful synchronization with the update channel:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-param-get param-name=last-update-sync uuid=$pool_uuid
```

Alternatively, you can manually synchronize your XenServer pool with the update channel:

```
1 pool_uuid=$(xe pool-list --minimal)
2 update_checksum=$(xe pool-sync-updates uuid=$pool_uuid --minimal)
```

The `update_checksum` unique identifier indicates the level of the updates installed. It changes whenever new updates are made available in the public CDN. This checksum is later used when applying updates to your pool to ensure that you are always applying the latest available updates. The value of `update_checksum` can also provide useful information if you need to contact Technical Support.

Note:

After synchronizing, apply the updates to your pool as soon as possible to benefit from the latest updates.

If you designate a new pool coordinator after synchronizing but before applying updates to the hosts in the pool, you must synchronize again with the new pool coordinator before you can update the pool.

Do not synchronize your XenServer pool while the pool is in the process of being updated.

Now that your pool is synchronized with the channel, you can view and apply the available updates. Next step: View available updates for your pool.

Offline updates

Offline updates have the following prerequisites:

- Your host must be installed with the ISO published on October 7, 2024, or later.
- You must use XenCenter version 2024.4.0 or later to apply update bundles.

Configure the offline updates channel

Follow these steps to configure your pool to use an offline updates channel:

1. Create the offline update channel for your pool:

```
1 xe repository-introduce-bundle name=label=<name-label> [name-  
description=<name-description>]
```

This command creates a bundle repository that acts as the Offline update channel and stores the update bundle for offline application. It returns the repository UUID. Note this UUID to use later in Install updates.

2. Enable the offline update channel for your pool:

```
1 xe pool-param-set repositories=<bundle-repository-uuid> uuid=<pool_uuid>
```

This command enables the bundle repository as the offline update channel for your pool.

Synchronize a new update bundle for your pool

After configuring your pool to use the offline update channel, you can follow these steps every time you want to synchronize and update from a new bundle.

1. Download the latest update bundle from the following page using your dedicated internet-connected machine: [Download XenServer update bundles](#).
2. Transfer the update bundle to the system where you run xe CLI commands. If you are transferring the bundle file onto a host in your pool, ensure that you delete it after synchronizing to reclaim disk space.
3. Upload and synchronize the update bundle file.

```
1 xe pool-sync-bundle filename=<bundle-file-path>
```

4. Delete the bundle file:

```
1 rm <bundle-file-path>
```

Note:

If you designate a new pool coordinator after synchronizing but before applying updates to the hosts in the pool, you must synchronize again with the new pool coordinator before you can update the pool.

Do not synchronize your XenServer pool while the pool is in the process of being updated.

Now that your pool is synchronized with the update bundle, you can view and apply the available updates. Next step: [View available updates for your pool](#).

Update a host to pool level

This method enables you to update a standalone XenServer 9 host to the same level of updates as a XenServer 9 pool. The pool must be at a more recent level of updates than the standalone host. Use

this process to more easily join a new host to an existing pool without having to update every host in the pool.

Pool-level update synchronizes the standalone host with the pool and updates from update files held on the pool coordinator. However, if the pool coordinator has changed since the pool was last updated, this capability might not be available for your pool because the current pool coordinator does not hold the update files.

Configure a pool as an update channel

Follow these steps to configure your host to use another pool as an updates channel:

1. On the coordinator of the pool you are using as the updates channel, get the server certificate:

```
1 xe host-get-server-certificate uuid=<remote_pool_coordinator_uuid>
```

The command outputs the certificate. Copy this information to your clipboard or keep it where you can access it in the next step.

2. On the host that you want to update, create a temporary file and copy into it the output from the previous command.

For example:

```
1 vi ~/tmp/certificate.pem
```

3. Check that the remote pool has the latest synchronized updates applied to it, before configuring it as an update channel. The following command must return **true**:

```
1 xe host-param-get uuid=<remote_pool_coordinator_uuid> param-name=
  latest-synced-updates-applied
```

If the command returns **false**, apply the synchronized updates to the remote pool before continuing to the next step.

4. Create the remote pool update channel for your host:

```
1 xe repository-introduce-remote-pool name-label=<name_label> binary
  -url=https://<url_to_pool_coordinator>/repository/enabled
  certificate-file=<path_to_certificate> [name-description=<name-
  description>]
```

- The value of the **binary-url** parameter is the URL address to the pool coordinator, either as an FQDN or IP address, with the path **/repository/enabled** appended.
- The value of the **certificate-file** parameter is the path to the temporary certificate file that you created in the previous step.

The command returns the repository UUID. Note the UUID for later use when installing updates.

1. Enable the remote pool update channel for your pool:

```
1 xe pool-param-set repositories=<remote_pool_repository_uuid>
```

This command enables the repository information on the remote pool coordinator as the update channel for your pool.

Synchronize the host with the pool

Manually synchronize your XenServer host with the remote pool update channel.

1. Run the following command on your standalone host:

```
1 xe pool-sync-updates username=root password=<password>
```

Provide the user name and password for the coordinator of the pool you are updating from.

The command returns a checksum that acts as a unique identifier indicating the level of the updates available to be installed.

Note:

After synchronizing, apply the updates to your pool as soon as possible to benefit from the latest updates.

Do not synchronize your XenServer host while the host is in the process of being updated.

Now that your host is synchronized with the pool, you can view and apply the available updates. Next step: View available updates for your pool.

View available updates for your pool

Before applying updates, view the available updates for your pool and complete any required pre-update tasks. For more information about the different update tasks, see Understand the guidance categories and update tasks.

Check for available updates for a particular host

To check if there are any available updates for a particular host after, run the following command:

```
1 xe host-param-get param-name=latest-synced-updates-applied uuid=<host_uuid>
```

This command returns **yes** if all the synchronized updates are applied to the host and **no** if there are synchronized updates available to be applied.

Note:

This command checks only synchronized updates. There might be newer updates available on the update channel that haven't yet been synchronized to the pool. For information about how to synchronize, see [Synchronize new updates for your pool for CDN updates](#) or [Synchronize a new update bundle for your pool for offline updates](#).

Check for available updates for all hosts

1. Check for available updates for all hosts by making a GET request on the HTTP endpoint [/updates](#). This retrieves the metadata for available updates. You can access the endpoint by sending a request to <https://<myserver>/updates>.

Use the [wget](#) utility:

```
1 wget -O - --no-check-certificate https://<user name>:<password>@<
   coordinator IP address>/updates
```

You can also use an HTTP client library with the following parameters:

```
1 HTTP GET
2 session_id: <XAPI session ID returned from login>
3 host_refs: <host XAPI reference>
```

2. Inspect the JSON response returned by the request. The response contains metadata about the available updates, including:
 - [hosts](#): Lists the available updates for individual hosts.
 - [updates](#): Lists the details of the available updates.
 - [hash](#): The [update_checksum](#) (used to ensure that you are always applying the latest available updates).
3. Review the [guidance](#) object within the [hosts](#) and [updates](#) objects. This object includes the following keys:
 - [mandatory](#)
 - [recommended](#)
 - [full](#)
 - [livepatch](#)

These categories outline all the update tasks associated with the updates that you are applying. Perform the necessary pre-update tasks based on the guidance categories. For more information about these categories, see [Understand the guidance categories and update tasks](#).

Next step: [Apply updates to your pool](#)

Understand the guidance categories and update tasks

Some tasks (such as evacuating or rebooting your hosts) might be required before and after applying updates to your pool. Sometimes, there are no update tasks required.

Guidance categories

XenServer tries to minimize the disruption to your VMs that these tasks might cause by categorizing the tasks into **Mandatory**, **Recommended**, **Full-effectiveness**, and **Live patch**. These categorizations enable you to judge whether an update task that might cause downtime or minor disruption for your hosts or VMs is necessary for your environment and risk profile.

Updates can have tasks listed in more than one of these categories. For example, an update might require that you restart the host to get the full-effectiveness of the update, but recommend restarting the toolstack to get most of the benefit of the update with less potential disruption to the pool.

During the update process, you can choose to carry out one of the following three levels of tasks:

1. Mandatory
2. Mandatory + Recommended
3. Mandatory + Recommended + Full-effectiveness

Mandatory Mandatory tasks *must* be performed after an update otherwise the system might fail at runtime. These actions are required to enable critical fixes and ensure that your environment is secure and stable. You cannot opt out of mandatory tasks.

Recommended Recommended tasks are the tasks that we recommend you perform to get the benefit from most of the features and fixes delivered in the updates. If you choose not to perform these tasks now, they are listed in the pending update tasks for the applicable pool, host, or VM.

Why perform the recommended tasks:

- These tasks are the ones that ensure a secure, stable XenServer environment.

Why opt out of the recommended tasks:

- After reviewing the detailed information for the updates, you judge that the risk of not fully applying these updates now is acceptable.
- The recommended tasks cause unwanted disruption to your VMs now.

If you have opted to defer the recommended tasks, plan to complete these tasks during a suitable maintenance window to maintain the stability of your environment.

Full-effectiveness Full effectiveness tasks are required to get the benefit of the related update. The updates that have full-effectiveness tasks associated with them are relevant only to users on certain hardware or using specific features.

Review the update information to understand whether these tasks are required for your environment. If you choose not to perform these tasks now, they are listed in the pending tasks for the applicable pool, host, or VM.

Why perform the full-effectiveness tasks:

- The updates that have full-effectiveness tasks are relevant to your hardware, environment, or configuration.

Why opt out of the full-effectiveness tasks:

- The updates that have full-effectiveness guidance are not relevant to your hardware, environment, or configuration.
- The full-effectiveness tasks cause unwanted disruption to your VMs now.
- You do not need the benefits of these updates right now.

If the full-effectiveness tasks apply to your environment, but you have opted to defer them, plan to complete these tasks during a suitable maintenance window to maintain the stability of your environment.

Live patches Updates to certain components can include a live patch. A live patch enables updates to be applied to core components without requiring a host reboot.

Whether a live patch can be applied to your hosts depends on the version of the component that was installed when the hosts were last rebooted. If an update can be applied as a live patch to your hosts, the live patch guidance replaces the recommended guidance.

Example:

You have two pools. Pool A is updated to a recent level. Pool B has not been updated for some time. We release a new update that has the recommended update task “Restart host” and the live patch update task “Restart toolstack”.

In pool A, the live patch can be applied to these more up-to-date hosts. The recommended guidance shows “Restart toolstack”. The less disruptive task from the live patch guidance overrides the recommended guidance.

In pool B, the live patch cannot be applied to the hosts as they are at an older level. The recommended guidance shows “Restart host”. The recommended guidance remains applicable. The live patch guidance is irrelevant in this case.

Sometimes only some of the fixes in an update are enabled when the update is applied as a live patch. In this situation, you might still see a recommended reboot for an update even though the live patch applies. Review the update details to understand whether you need all fixes in the update or just those fixes enabled by the live patch. You can then use this information to choose whether to perform the tasks. For more information, see [View available updates for your pool](#).

Update tasks

Applying updates might require you to complete one or more tasks, either before or after the updates are applied. These tasks can appear in any guidance category. The following tables outline these tasks, their associated xe CLI commands, and descriptions of what they involve.

Update tasks for hosts Tasks required *before* applying updates:

Update task	xe CLI command to carry out task	Description
Evacuate host	<code>xe host-evacuate</code>	All VMs must be migrated off the XenServer host or shutdown before applying the update. To complete this task, XenCenter migrates any VMs off the host. While this task is in progress, the XenServer pool is operating at reduced capacity as one host is temporarily unavailable to run VMs. This task is sometimes carried out as part of the 'Reboot host' task.

Tasks required *after* applying updates:

Update task	xe CLI command to carry out task	Description
Reboot host	<code>xe host-reboot</code>	The XenServer host must be restarted. Any VMs are migrated off the host and the host is restarted. While this task is in progress, the XenServer pool is operating at reduced capacity as one host is temporarily unavailable to run VMs.
Reboot host on Xen live patch failure	<code>xe host-reboot</code>	Applying a Xen live patch failed. The XenServer host must be restarted to get the update to take effect. Any VMs are migrated off the host and the host is restarted. While this task is in progress, the XenServer pool is operating at reduced capacity as one host is temporarily unavailable to run VMs.
Reboot host on kernel live patch failure	<code>xe host-reboot</code>	Applying a dom0 kernel live patch failed. The XenServer host must be restarted to get the update to take effect. Any VMs are migrated off the host and the host is restarted. While this task is in progress, the XenServer pool is operating at reduced capacity as one host is temporarily unavailable to run VMs.

Update task	xe CLI command to carry out task	Description
Restart toolstack	<code>xe-toolstack-restart</code>	The toolstack on the host must be restarted. If performed on the pool coordinator, the connection to the pool is temporarily lost. Restarting on other hosts has no visible effect on the pool's state.

Update tasks for VMs Some updates introduce new features for VMs, requiring the following tasks:

Update task	xe CLI command to carry out task	Description
Restart VM	<code>xe vm-reboot</code>	The VM must be restarted. In XenCenter, the VM displays a red stop icon (square on red) while restarting. Once the task is complete, the green play icon appears, indicating the VM is running. During the reboot, the VM is unavailable to end users.
Restart device model	<code>xe vm-restart-device-models</code>	The device model for VMs on the updated host must be restarted. In XenCenter, the VM shows a yellow play triangle while the device model restarts. Once complete, the green play icon appears. During this time, you cannot stop, start, or migrate the VM. The end user of the VM might experience a brief pause and resume in their session. To support this task on a Windows VM, the VM must have XenServer VM Tools for Windows installed.

Apply updates to your pool

Before you start

- Ensure that all the hosts in your pool are online before carrying out the pool update.
- Ensure that there are no pending mandatory update tasks on any host or VM. Any mandatory update tasks pending from previous updates must be carried out before starting a new pool update. For more information, see [View the update tasks required for your host](#) and [View the update tasks required for your VM](#).
- Disable High Availability (HA) if it is enabled:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-ha-disable uuid=$pool_uuid
```

- Disable Workload Balancing (WLB) if it is enabled:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-param-set wlb-enabled=false uuid=$pool_uuid
```

Install updates

To perform an update to your pool, you must apply updates on every host in the pool, starting with the pool coordinator first. Follow these steps, starting with the pool coordinator:

1. Disable the host:

```
1 xe host-disable uuid=<host_uuid>
```

2. If one of the update tasks required for the update is 'Evacuate host' or 'Reboot host', evacuate the host:

```
1 xe host-evacuate uuid=<host_uuid>
```

If you can't migrate a VM to other hosts during the host evacuation, shut down or suspend the VM.

3. Apply updates to the host:

- a) Get the update checksum of the updates to apply:

```
1 xe repository-param-get uuid=<repository UUID> param-name=hash
```

- b) Use this checksum to apply the set of updates:

```
1 xe host-apply-updates uuid=<host_uuid> hash=<update_checksum>
```


4. Get a list of the host update tasks required. For more information, see [Update tasks for your host](#).

Carry out the host's update tasks in the list in the following order:

- a) Restart toolstack (can be skipped if there is a 'Reboot host' to be carried out)
 - b) Reboot host
5. For every running VM on the host, get a list of the VM update tasks required. For more information, see [Update tasks for your VM](#).

Carry out the VM's update tasks in the list in the following order:

- a) Restart device model (can be skipped if there is a 'Restart VM' to be carried out)
 - b) Restart VM
6. Enable the host if it is still in a disabled state:

```
1 xe host-enable uuid=<host_uuid>
```

7. For every VM which you migrated to another host using [host-evacuate](#) before the host update, get a list of the update tasks. For more information, see [Update tasks for your VM](#).

If 'Restart VM' is in the list of update tasks, shut down the VM and start it on the current updated host. Otherwise, migrate the VM back to the current updated host.

8. Resume or start the VMs that you shut down or suspended before you applied updates.
9. View the host update status:

```
1 xe host-param-get param-name=last-software-update uuid=<host_uuid>
2 xe host-param-get param-name=latest-synced-updates-applied uuid=<host_uuid>
3 xe host-param-get param-name=last-update-hash uuid=<host_uuid>
```

Repeat the preceding steps to update every host in your pool.

After updating your hosts

After updating each host in your pool, carry out any remaining update tasks.

1. For every VM in your pool, get a list of the update tasks. For more information, see [Update tasks for your VM](#).
If 'Restart VM' is in the list of update tasks, carry it out.
2. Enable HA if you disabled it before applying updates:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-ha-enable uuid=$pool_uuid
```

3. Enable WLB if you disabled it before applying updates:

```
1 pool_uuid=$(xe pool-list --minimal)
2 xe pool-param-set wlb-enabled=true uuid=$pool_uuid
```

4. If you chose to carry out only the mandatory update tasks required for a pool update, the update tasks that have not been carried out are appended to the list of pending update tasks required for your hosts. To view this list and carry out these tasks, see [Complete any pending update tasks](#).

Complete any pending update tasks

After applying updates, complete any pending update tasks to ensure your hosts and VMs operate as expected. Use the following commands to view the required tasks.

View pending tasks for a host

Get a list of the mandatory tasks for your host:

```
1 xe host-param-get param-name=pending-guidances uuid=<host_uuid>
```

Get a list of the recommended tasks for your host:

```
1 xe host-param-get param-name=pending-guidances-recommended uuid=<
  host_uuid>
```

Get a list of the full-effectiveness tasks for your host:

```
1 xe host-param-get param-name=pending-guidances-full uuid=<host_uuid>
```

View pending tasks for a VM

Get a list of the mandatory tasks for your VM:

```
1 xe vm-param-get param-name=pending-guidances uuid=<VM UUID>
```

Get a list of the recommended tasks for your VM:

```
1 xe vm-param-get param-name=pending-guidances-recommended uuid=<VM UUID>
```

Get a list of the full-effectiveness tasks for your VM:

```
1 xe vm-param-get param-name=pending-guidances-full uuid=<VM UUID>
```

XenCenter®

December 17, 2025

With the XenCenter management console, you can manage your XenServer® environment and deploy, manage, and monitor virtual machines from an application on your Windows desktop machine.

DOWNLOAD XENCENTER

XenCenter documentation

The XenCenter management console has its own set of documentation at [XenCenter product documentation](#).

Getting started

Let's get started with the XenCenter management console:

1. You can install XenCenter on most Windows systems, as long as your system has .NET 4.8 or later installed. For more information, see [System requirements](#).
2. To install XenCenter, run the installer file you downloaded from [our website](#). The installer guides you through the installation process. For more information, see [Install XenCenter](#).
3. Now that you have XenCenter installed on your Windows system, take a moment to get familiar with the layout of the UI and its capabilities. For more information, see [Exploring the XenCenter workspace](#).
4. Connect XenCenter to your XenServer hosts and start to manage them. For more information, see [Add New Server](#).

What's new in XenCenter

We release updates to XenCenter frequently containing new features, improvements, and bug fixes. To find out what's new in the latest version of XenCenter, see [What's New](#).

Using XenServer® with Citrix® products

June 3, 2026

XenServer provides features that enhance its interoperation with Citrix Virtual Apps and Desktops™, Citrix DaaS, and Citrix Provisioning.

For more information about these products, see:

- [Citrix Virtual Apps and Desktops product documentation](#)
- [Citrix DaaS product documentation](#)
- [Citrix Provisioning product documentation](#)

Supported versions

You can find the versions of these products that XenServer 8.4 interoperates with on the Citrix website: [Supported Hypervisors for Citrix Virtual Apps and Desktops \(MCS\) and Citrix Provisioning \(PVS\)](#).

Licensing

To use XenServer for your Citrix Virtual Apps and Desktops or Citrix DaaS workloads, you must have a XenServer Premium Edition license. For more information, see [Licensing](#).

XenServer is an entitlement of the Citrix for Private Cloud, Citrix Universal Hybrid Multi-Cloud, Citrix Universal MSP, and Citrix Platform License subscriptions for running your Citrix workloads. [Read more](#).

You can find more information about XenServer licenses at the [XenServer website](#).

XenServer features for Citrix Virtual Apps and Desktops

The following XenServer features are designed for use with Citrix Virtual Apps and Desktops, Citrix DaaS, and Citrix Provisioning™:

- **XenCenter Citrix integration:** Connect XenCenter to Citrix Virtual Apps and Desktops to view information about your XenServer VMs and their place in your Citrix environment.
- **IntelliCache:** Using XenServer with IntelliCache makes hosted Citrix Virtual Desktop deployments more cost-effective by enabling you to use a combination of shared storage and local storage. It is of particular benefit when many VMs all share a common OS image. The load on the storage array is reduced and performance is enhanced. In addition, network traffic to and

from shared storage is reduced as the local storage caches the primary image from shared storage.

- **Read caching:** Read caching improves a VM's disk performance by caching data within the host's free memory. It improves performance in a Citrix Virtual Desktops Machine Creation Services (MCS) environment where many VMs are cloned off a single base VM, as it drastically reduces the number of blocks read from disk.
- **PVS-Accelerator:** The XenServer PVS-Accelerator feature offers extended capabilities for customers using XenServer with Citrix Provisioning. PVS-Accelerator provides many benefits including data locality, improved end-user experience, accelerated VM boots and boot storms, simplified scale-out by adding more hypervisor hosts, and reduced TCO and simplified infrastructure requirements.
- **Smooth roaming support for Virtual Desktop Tablet Mode:** XenServer, in conjunction with Citrix Virtual Apps and Desktops, enables you to experience Windows 10 Continuum experience in a virtualized environment.
- **Graphics virtualization:** XenServer provides the virtual delivery of 3D professional graphics applications and workstations in XenServer through GPU Pass-through (for NVIDIA, AMD, and Intel GPUs) and hardware-based GPU sharing with NVIDIA vGPU.

Migrating from other hypervisors

By using the XenCenter® **Import** wizard, Citrix customers can easily migrate their VMs from VMware to XenServer.

Learn more about moving from VMware to XenServer in our Tech Zone article: [VMware to XenServer migration guide](#).

Migrating from VMware

As part of its software streaming technology, Citrix Provisioning stores a shared disk image (vDisk) as a VHDX/AVHDX file.

VHDX is a virtual hard disk format that is used to store the disk data for VMs. It is an improved version of the VHD format that provides better performance, fault tolerance, and data protection features in addition to a larger capacity.

AVHDX (Automatic Virtual Hard Disk) is a disk image format that is part of the VHDX standard for virtual hard disks. It is used to store the snapshot or checkpoint information of VMs. Whenever you create a snapshot on a VM, an AVHDX file is automatically generated to preserve the current state of the virtual hard disk. This file can be used to track the snapshot information of VMs. Each AVHDX file is linked to

the previous AVHDX file, creating a chain that contains all changes made since the base virtual hard disk file.

By importing a VHDX/AVHDX file using the XenCenter **Import** wizard, you can migrate your vDisk from VMware to XenServer. For information on how to use the XenCenter **Import** wizard to import a VHDX/AVHDX file, see [Import disk images](#).

Limitations You cannot import a VHDX/AVHDX file that is larger than 2 TB.

Best practices

When configuring and managing your XenServer environment there are steps you can take to optimize how it works with Citrix products.

For in-depth articles that describe reference architectures and deployments, see our [TechZone](#).

Installing and upgrading

- When first installing XenServer hosts, you can enable Intellicache to cache VM data locally and improve performance. For more information, see [Intellicache](#).
- If you are upgrading from an earlier version of XenServer, the method you use for this upgrade can depend on your Citrix Virtual Apps and Desktops workload. For more information, see [Upgrade scenarios for Citrix Virtual Apps and Desktops](#).

Configuring your environment

- The XenServer host comes installed with a default TLS certificate. However, to use HTTPS to secure communication between XenServer and Citrix Virtual Apps and Desktops, install a certificate provided by a trusted certificate authority. For more information, see [Install a TLS certificate on your host](#).

Memory usage

- When XenServer is first installed, it allocates a certain amount of memory to the control domain. In many Citrix Virtual Apps and Desktops environments, it is advisable to increase the amount of memory allocated to the control domain beyond this default.

Increase the control domain memory in the following cases:

- You are running many VMs on the server

- You are using PVS-Accelerator
- You are using read caching

For information about changing the amount of control domain memory and monitoring the memory behavior, see [Memory usage](#).

[This is an embedded video. Click the link to watch the video](#)

Upgrade scenarios for XenServer and Citrix Virtual Apps and Desktops™

March 18, 2026

XenServer contains features and optimizations that make it an ideal hypervisor to use in your Citrix Virtual Apps and Desktops environment.

If you are using XenServer with Citrix Virtual Apps and Desktops, there are some considerations when performing your upgrade that are not covered in the main upgrade article: [Upgrade from an existing version](#). Review both this article and the main upgrade article before starting your upgrade from XenServer 8.4 to XenServer 9.

Considerations when upgrading XenServer® in a Citrix Virtual Apps and Desktops environment:

- XenServer hosts restart twice as part of an upgrade. At the beginning of the upgrade, you must boot your server into the installation media. At the end of the process, the installer restarts the server to complete the upgrade. VMs on these hosts must be either migrated or stopped during this time.
- The approach to use for upgrading XenServer depends on your XenServer environment, your Citrix Virtual Apps and Desktops environment, and the types of machines and applications being hosted by XenServer.
- You might be required to do some preparation in your Citrix Virtual Apps and Desktops environment before starting your XenServer upgrade.
- This article only covers use cases where the Citrix Virtual Apps and Desktops workload is hosted in the XenServer pool. Cases where you are also hosting parts of your Citrix Virtual Apps and Desktops infrastructure on VMs in the XenServer pool are not covered by this article. Factor these components in when doing your upgrade planning.
- Ensure that the version of Citrix Virtual Apps and Desktops you are using is supported for both the version of XenServer you are upgrading from and the version you are upgrading to. For more information, see [Supported Hypervisors for Citrix Virtual Apps and Desktops \(MCS\) and Citrix Provisioning \(PVS\)](#).
- The time it takes to do the upgrade and the potential for service outage depends on your upgrade approach. The full upgrade of an entire pool might take multiple hours to complete.

- This article assumes that the time to fully upgrade a single XenServer host is 35 minutes. This host upgrade time includes the upgrade process and any required restarts.

The approaches described in this article aim to guide you to an upgrade method that reduces the possibility of service outage and enables the upgrade process to fit within your maintenance window. However, in some cases, service outages are unavoidable. If the XenServer upgrade process cannot fit within your maintenance window, you can run your pool in mixed mode for a short time between maintenance windows. However, this is not recommended. For more information, see [Mixed-mode pools](#).

During your planned XenServer upgrade maintenance window, follow these restrictions:

- Do not attempt to reconfigure the infrastructure of the pool that is being upgraded. For example, do not add or eject hosts from the pool.
- Do not add, start, or stop any VMs in the pool that is being upgraded.
- Do not perform catalog updates during the window.

Rolling Pool Upgrade

Rolling Pool Upgrade is a XenServer feature designed to make the upgrade process easier and to minimize downtime.

The **Rolling Pool Upgrade** wizard in XenCenter® guides you through the upgrade procedure and organizes the upgrade path automatically. For pools, each of the servers in the pool is upgraded in turn, starting with the pool coordinator. Before starting an upgrade, the wizard conducts a series of pre-checks. These pre-checks ensure certain pool-wide features, such as high availability, are temporarily disabled and that each server in the pool is prepared for upgrade. Only one server is offline at a time. Any running VMs are automatically migrated off each server before the upgrade is installed on that server.

You can use Rolling Pool Upgrade for many of the Citrix Virtual Apps and Desktops use cases outlined in this article. For each, the upgrade time is the same: the number of hosts in the pool multiplied by the upgrade time for a single host. (**N x 35 minutes**). The potential for VM outage depends on your Citrix Virtual Apps and Desktops workload and XenServer pool setup.

Even if you intend to use Rolling Pool Upgrade to upgrade your XenServer pool, review the information for your specific environment to ensure that you understand the Citrix Virtual Apps and Desktops prerequisite actions, any special considerations, and the behavior to expect.

Use cases

This article identifies several broad use cases. For each of these use cases, we assume that the XenServer pool hosts only one type of Citrix Virtual Apps and Desktops workload. If your pool

contains a mix of different types of workload, review all the cases that apply to your pool to decide what your preferred upgrade approach is.

First, consider how your XenServer environment is configured:

- XenServer pool with shared storage

In a XenServer pool with one or more shared storage repositories (SRs), the VM disks can be hosted on this shared storage, which enables the VMs to migrate between hosts during the upgrade. This configuration can reduce or remove the need for VM downtime.

- XenServer pool without shared storage or a standalone host

In a XenServer pool without shared storage or on a standalone XenServer host, the VMs can't migrate during the upgrade process. When the host reboots as part of the upgrade, you must shut down the VMs.

XenServer pool with shared storage

If you are upgrading a pool where the VM disks are located on shared storage, you can evacuate VMs from each XenServer host in the pool while it is upgraded.

Most use cases on this type of pool can be upgraded by using Rolling Pool Upgrade. However, the prerequisite actions required in Citrix Virtual Apps and Desktops and the outage behavior is different depending on your workload.

Consider what type of Citrix Virtual Apps and Desktops workload is hosted in your pool:

- Unassigned single-session desktops
- Other workloads

XenServer pool without shared storage or a standalone host

If you are upgrading a pool where the VM disks are located on local storage or you have a single host in your pool, the VMs cannot be migrated off the XenServer hosts while they are upgraded. In these cases, the VMs must be shut down for the duration of the host or pool upgrade. Some outage to your virtual apps and desktops is unavoidable in these cases.

Consider what type of Citrix Virtual Apps and Desktops workload is hosted in your pool:

- Assigned desktops
- Other workloads

Case 1: Single-session desktops running on a pool with shared storage

This use case covers XenServer pools with shared storage whose primary workload is single-session virtual desktops with the random machine allocation type. Machines of this type must be managed by either Citrix Provisioning or by Machine Creation Services.

For any workload that is managed by Citrix Virtual Apps and Desktops, including those that are power managed by Citrix Provisioning™ and Machine Creation Services, you cannot maintain a complete workload while the upgrade is being performed. Power management of machines can be problematic during the upgrade process and you cannot disable power management without also disabling new session creation.

Recommended options for upgrade:

- Rolling Pool Upgrade
 - Estimated upgrade time: The number of hosts in the pool multiplied by the upgrade time for a single host. **(N x 35 minutes)**
 - Outage behavior: All machines are in Citrix Virtual Apps and Desktops maintenance mode for the entire upgrade time.

If possible, make the workload available from other XenServer pools with capacity during the upgrade of this pool. This approach might cause reduced capacity during the upgrade. If you do not have capacity for the workload on your other XenServer hosts and pools, we recommend that you declare an outage for all machines in your workload.

Rolling Pool Upgrade (1)

Review the steps and guidance in [Before you start](#).

1. Put all the machines in the pool in maintenance mode. If all of the machines are using the same connection, you can put the entire machine catalog in maintenance mode.
2. Notify all impacted users of the impending outage.
 - If sessions are still running on the machines in this pool, ask users to log off or force their sessions to end.
 - Inform users that after they log off they can't log in again until full service is resumed.
3. In XenCenter, start the **Rolling Pool Upgrade** wizard and choose automatic mode. For more information, see [Rolling Pool Upgrade by using XenCenter](#).

When the upgrade is complete, any VMs that were suspended as part of the **Rolling Pool Upgrade** are restarted.

4. Take the machines out of maintenance mode.

New sessions can now be started and full service resumed.

Case 2: Other workloads running on a pool with shared storage

This use case covers XenServer pools with shared storage whose primary workload is either single-session virtual desktops with the assigned machine allocation type or multiple-session virtual applications with the random machine allocation type.

Recommended options for upgrade:

- Rolling Pool Upgrade
 - Estimated upgrade time: The number of hosts in the pool multiplied by the upgrade time for a single host. (**N x 35 minutes**)
 - Outage behavior: No service outage

Rolling Pool Upgrade (2)

Review the steps and guidance in [Before you start](#).

1. Ensure that the pool has enough capacity to run your workload with one fewer host in the pool. During the upgrade process, each host is removed one at a time. The remaining hosts must be able to run all required VMs.

If there is not enough capacity in the pool, some machines might not be available during the upgrade process. If possible, you can suspend any non-critical VMs during the upgrade process.

2. Ensure that all machines provided by the XenServer pool are powered on and are registered with Citrix Virtual Apps and Desktops™ in the relevant Delivery Groups.
 - For unmanaged machines:
 - Use XenCenter® to confirm that all VMs are powered on.
 - Do not perform any manual power actions during the upgrade process.
 - For power-managed machines:
 - Ensure that all machines are powered on (by using XenCenter, Citrix® Studio, or Web Studio).
 - **To enable new sessions to start during the upgrade process:**
 - Do not set the machines in maintenance mode.
 - Do not perform any manual power actions during the upgrade process.
 - Disable any power management schemes that might suspend machines.

- Ensure that there are no other processes that might power off or suspend the machines.
- **If it's acceptable for new sessions to be unable to start during the upgrade:**
 - Put the hosting connection in maintenance mode. For more information, see [Turn maintenance mode on or off for a connection](#).
 - Inform end users that if they log off they cannot reconnect for the time of the upgrade.

For more information, see [Power managed machines in a delivery group](#).

- For machines managed by Machine Creation Services™
 - Follow the same guidance as for power-managed machines in the preceding list item.
 - In addition, do not attempt to create new machines during the entire upgrade period.
3. In XenCenter, start the **Rolling Pool Upgrade** wizard and choose automatic mode. For more information, see [Rolling Pool Upgrade by using XenCenter](#).
 4. Restore your environment operations to their usual configuration.
 - Remove any maintenance mode flags set in earlier steps.
 - Revert any power management scheme adjustments made in the earlier steps.

Case 3: Assigned desktops running on a pool with local storage or on a standalone host

This use case covers XenServer standalone hosts or pools that do not have shared storage whose primary workload is either single-session virtual desktops with the assigned machine allocation type.

Recommended options for upgrade:

- Rolling Pool Upgrade Use RPU in automatic mode in a single maintenance window. This requires all users to have outage for the entire upgrade, but has a lower administration overhead for a pool
 - Estimated upgrade time: The number of hosts in the pool multiplied by the upgrade time for a single host. **(N x 35 minutes)**
 - Outage behavior: All machines are in Citrix Virtual Apps and Desktops maintenance mode for the entire upgrade time.
- Manual upgrade This mode provides the least outage for each user during the upgrade, but is more involved for the administrator
 - Estimated upgrade time: Two times the upgrade time for a single host. **(Approximately 70 minutes)**
 - Outage behavior: Each desktop is unavailable during the upgrade time for their individual host. This time is typically 35 minutes.

Rolling Pool Upgrade (3)

Review the steps and guidance in [Before you start](#).

1. Put all delivery groups or catalogs that are providing machines from the pool into maintenance mode.

While the machines are in maintenance mode, new sessions cannot be started on machines in the pool. Existing sessions are maintained until the machines are shut down or suspended.

For more information, see [Prevent users from connecting to a machine in a delivery group](#).

2. Notify all affected users of the impending outage. Provide a time by which they must end their sessions and indicate when service will be restored.
3. Check for remaining sessions on impacted machines and take appropriate actions for these sessions.
4. In XenCenter, start the **Rolling Pool Upgrade** wizard and choose automatic mode. For more information, see [Rolling Pool Upgrade by using XenCenter](#).

When the upgrade is complete, any VMs that were suspended as part of the **Rolling Pool Upgrade** are restarted.

5. Take the machines out of maintenance mode.

New sessions can now be started and full service resumed.

Manual upgrade (3)

You can use this manual process to upgrade the pool coordinator first then all other hosts in parallel to reduce overall outage time significantly.

Note:

With the parallel upgrade approach, the risk profile changes. If there is an issue during the upgrade, it might not be detected until all hosts have been upgraded and are experiencing the issue. Whereas if you upgrade your hosts sequentially, you can verify that the upgrade is successful on each host before going on to the next.

Review the steps and guidance in [Before you start](#).

1. Ensure that all machines provided by the XenServer pool or host are turned on and are registered with Citrix Virtual Apps and Desktops™ in the relevant delivery groups.
 - For unmanaged machines:
 - Use XenCenter® to confirm that all VMs are powered on.

- Do not perform any manual power actions during the upgrade process.
- For power-managed machines:
 - Ensure that all machines are powered on (using XenCenter or Studio).
 - **To enable new sessions to start during the upgrade process:**
 - Do not set the machines in maintenance mode.
 - Do not perform any manual power actions during the upgrade process.
 - Disable any power management schemes that might suspend machines.
 - Ensure that there are no other processes that might power off or suspend the machines.
 - **If it's acceptable for new sessions to be unable to start during the upgrade:**
 - Put the hosting connection in maintenance mode. For more information, see [Turn maintenance mode on or off for a connection](#).
 - Inform end users that if they log off they cannot reconnect for the time of the upgrade.

For more information, see [Power managed machines in a delivery group](#).

- For machines managed by Machine Creation Services™
 - Follow the same guidance as for power-managed machines in the preceding list item.
 - In addition, do not attempt to create machines during the entire upgrade period.
2. Identify the pool coordinator and associated VMs.
 3. Put the machines in the catalog on the pool coordinator host into maintenance mode.
 4. Use Director, Citrix® Studio, or Web Studio to send messages to users that are still connected to active sessions, warning them that their desktop is going offline for a period. This period is the upgrade time for this individual host (approximately 35 minutes).
 5. Update the pool coordinator by using the xe CLI:
 - a) Disable the pool coordinator. This prevents any new VMs from starting on or being migrated to the specified host.

```
1 xe host-disable host=<uuid_or_name_label>
```

- b) Ensure that no VMs are running on the pool coordinator. Shut down, suspend, or migrate VMs to other hosts in the pool.

- To shut down a VM, use the following command:

```
1 xe vm-shutdown
```

- To suspend a VM, use the following command:

```
1 xe vm-suspend
```

- To migrate a specific VM, use the following command:

```
1 xe vm-migrate
```

Migrating specified VMs to specified hosts gives you full control over the distribution of migrated VMs to other hosts in the pool.

- To evacuate the host, use the following command:

```
1 xe host-evacuate
```

Evacuating all VMs from a host leaves the distribution of migrated VMs to XenServer.

- c) Shut down the pool coordinator.

```
1 xe host-shutdown
```

Important:

You are unable to contact the pool coordinator until the upgrade of the pool coordinator is complete. Shutting down the pool coordinator causes the other hosts in the pool to enter *emergency mode*. Hosts can enter emergency mode when they in a pool whose pool coordinator has disappeared from the network and cannot be contacted after several attempts. VMs continue to run on hosts in emergency mode, but control operations are not available.

- d) Boot the pool coordinator using the XenServer installation media and method of your choice (such as, USB or network).
- e) Follow the XenServer installation procedure until the installer offers you the option to upgrade. Choose to upgrade.

When your pool coordinator restarts, the other hosts in the pool leave emergency mode and normal service is restored after a few minutes.

- f) Start or resume any shutdown or suspended VMs.
- g) Migrate any VMs that you want back to the pool coordinator.

If anything interrupts the upgrade of the pool coordinator or if the upgrade fails for any reason, do not attempt to proceed with the upgrade. Reboot the pool coordinator and restore to a working version.

6. After the pool coordinator is upgraded, take the machines on the pool coordinator out of maintenance mode in Citrix Studio or Web Studio.
7. Complete the following steps in parallel for all remaining hosts in the pool:
- a) Put the machines in the catalog on the host into maintenance mode.

- b) Use Director, Citrix Studio, or Web Studio to send messages to users that are still connected to active sessions, warning them that their desktop is going offline for a period. This period is the upgrade time for this individual host (approximately 35 minutes).
- c) Disable the host by using the xe CLI.

```
1 xe host-disable host-selector=<host_selector_value>
```

- d) Ensure that no VMs are running on the host. Shut down, suspend, or migrate VMs to other hosts in the pool.

- To shut down a VM, use the following command:

```
1 xe vm-shutdown
```

- To suspend a VM, use the following command:

```
1 xe vm-suspend
```

- To migrate a specific VM, use the following command:

```
1 xe vm-migrate
```

Migrating specified VMs to specified hosts gives you full control over the distribution of migrated VMs to other hosts in the pool.

- To evacuate the host, use the following command:

```
1 xe host-evacuate
```

Evacuating all VMs from a host leaves the distribution of migrated VMs to XenServer.

- e) Shut down the host.

```
1 xe host-shutdown
```

- f) Boot the host using the XenServer installation media and method of your choice (such as, USB or network).
- g) Follow the XenServer installation procedure until the installer offers you the option to upgrade. Choose to upgrade.
- h) After the host upgrade is complete, start or resume any shutdown or suspended VMs.
- i) Migrate any VMs that you want back to the host.

If the upgrade of a subordinate host fails or is interrupted, you do not have to revert. Run the command `xe host-forget` in the pool to forget that host. Reinstall XenServer on the host, and then join it, as a new host, to the pool using the command `xe pool-join`.

- 8. After the XenServer hosts are updated, take the machines out of maintenance mode in Citrix Studio or Web Studio.

Case 4: Other workloads running on a pool with local storage or on a standalone host

This use case covers XenServer pools with shared storage whose primary workload is single-session virtual desktops or multiple-session virtual applications with the random machine allocation type.

For any workload that is managed by Citrix Virtual Apps and Desktops, including those that are power managed by Citrix Provisioning and Machine Creation Services, you cannot maintain a complete workload while the upgrade is being performed. Power management of machines can be problematic during the upgrade process and you cannot disable power management without also disabling new session creation.

Recommended options for upgrade:

- Rolling Pool Upgrade
 - Estimated upgrade time: The number of hosts in the pool multiplied by the upgrade time for a single host. (**N x 35 minutes**)
 - Outage behavior: All machines are in Citrix Virtual Apps and Desktops maintenance mode for the entire upgrade time.
- Manual upgrade
 - Estimated upgrade time: Two times the upgrade time for a single host. (**Approximately 70 minutes**)
 - Outage behavior: All machines are in Citrix Virtual Apps and Desktops maintenance mode for the entire upgrade time.

If possible, make the workload available from other XenServer pools with capacity during the upgrade of this pool. This approach might cause reduced capacity during the upgrade. If you do not have capacity for the workload on your other XenServer hosts and pools, we recommend that you declare an outage for all machines in your workload.

Rolling Pool Upgrade (4)

Review the steps and guidance in [Before you start](#).

1. Put all the machines in the pool in maintenance mode. If all of the machines are using the same connection, you can put the entire machine catalog in maintenance mode.
2. Notify all impacted users of the impending outage.
 - If sessions are still running on the machines in this pool, ask users to log off or force their sessions to end.
 - Inform users that after they log off they can't log in again until full service is resumed.

3. In XenCenter, start the **Rolling Pool Upgrade** wizard and choose automatic mode. For more information, see [Rolling Pool Upgrade by using XenCenter](#).

When the upgrade is complete, any VMs that were suspended as part of the **Rolling Pool Upgrade** are restarted.

4. Take the machines out of maintenance mode.

New sessions can now be started and full service resumed.

Manual upgrade (4)

You can use this manual process to upgrade the pool coordinator first then all other hosts in parallel to reduce overall outage time significantly.

Note:

With the parallel upgrade approach, the risk profile changes. If there is an issue during the upgrade, it might not be detected until all hosts have been upgraded and are experiencing the issue. Whereas if you upgrade your hosts sequentially, you can verify that the upgrade is successful on each host before going on to the next.

Review the steps and guidance in [Before you start](#).

1. Ensure that all machines provided by the XenServer pool or host are turned on and are registered with Citrix Virtual Apps and Desktops™ in the relevant delivery groups.
 - For unmanaged machines:
 - Use XenCenter® to confirm that all VMs are powered on.
 - Do not perform any manual power actions during the upgrade process.
 - For power-managed machines:
 - Ensure that all machines are powered on (using XenCenter or Studio).
 - **To enable new sessions to start during the upgrade process:**
 - Do not set the machines in maintenance mode.
 - Do not perform any manual power actions during the upgrade process.
 - Disable any power management schemes that might suspend machines.
 - Ensure that there are no other processes that might power off or suspend the machines.
 - **If it's acceptable for new sessions to be unable to start during the upgrade:**
 - Put the hosting connection in maintenance mode. For more information, see [Turn maintenance mode on or off for a connection](#).
 - Inform end users that if they log off they cannot reconnect for the time of the upgrade.

For more information, see [Power managed machines in a delivery group](#).

- For machines managed by Machine Creation Services™
 - Follow the same guidance as for power-managed machines in the preceding list item.
 - In addition, do not attempt to create machines during the entire upgrade period.
- 2. Identify the pool coordinator and associated VMs.
- 3. Put the machines in the catalog on the pool coordinator host into maintenance mode.
- 4. Use Director, Citrix® Studio, or Web Studio to send messages to users that are still connected to active sessions, warning them that their desktop is going offline for a period. This period is the upgrade time for this individual host (approximately 35 minutes).
- 5. Update the pool coordinator by using the xe CLI:

- a) Disable the pool coordinator. This prevents any new VMs from starting on or being migrated to the specified host.

```
1 xe host-disable host=<uuid_or_name_label>
```

- b) Ensure that no VMs are running on the pool coordinator. Shut down, suspend, or migrate VMs to other hosts in the pool.

- To shut down a VM, use the following command:

```
1 xe vm-shutdown
```

- To suspend a VM, use the following command:

```
1 xe vm-suspend
```

- To migrate a specific VM, use the following command:

```
1 xe vm-migrate
```

Migrating specified VMs to specified hosts gives you full control over the distribution of migrated VMs to other hosts in the pool.

- To evacuate the host, use the following command:

```
1 xe host-evacuate
```

Evacuating all VMs from a host leaves the distribution of migrated VMs to XenServer.

- c) Shut down the pool coordinator.

```
1 xe host-shutdown
```

Important:

You are unable to contact the pool coordinator until the upgrade of the pool coordinator is complete. Shutting down the pool coordinator causes the other hosts in the pool to enter *emergency mode*. Hosts can enter emergency mode when they in a pool whose pool coordinator has disappeared from the network and cannot be contacted after several attempts. VMs continue to run on hosts in emergency mode, but control operations are not available.

- d) Boot the pool coordinator using the XenServer installation media and method of your choice (such as, USB or network).
- e) Follow the XenServer installation procedure until the installer offers you the option to upgrade. Choose to upgrade.

When your pool coordinator restarts, the other hosts in the pool leave emergency mode and normal service is restored after a few minutes.

- f) Start or resume any shutdown or suspended VMs.
- g) Migrate any VMs that you want back to the pool coordinator.

If anything interrupts the upgrade of the pool coordinator or if the upgrade fails for any reason, do not attempt to proceed with the upgrade. Reboot the pool coordinator and restore to a working version.

6. After the pool coordinator is upgraded, take the machines on the pool coordinator out of maintenance mode in Citrix Studio or Web Studio.
7. Complete the following steps in parallel for all remaining hosts in the pool:
 - a) Put the machines in the catalog on the host into maintenance mode.
 - b) Use Director, Citrix Studio, or Web Studio to send messages to users that are still connected to active sessions, warning them that their desktop is going offline for a period. This period is the upgrade time for this individual host (approximately 35 minutes).
 - c) Disable the host by using the xe CLI.

```
1 xe host-disable host-selector=<host_selector_value>
```

- d) Ensure that no VMs are running on the host. Shut down, suspend, or migrate VMs to other hosts in the pool.
 - To shut down a VM, use the following command:

```
1 xe vm-shutdown
```
 - To suspend a VM, use the following command:

```
1 xe vm-suspend
```

- To migrate a specific VM, use the following command:

```
1 xe vm-migrate
```

Migrating specified VMs to specified hosts gives you full control over the distribution of migrated VMs to other hosts in the pool.

- To evacuate the host, use the following command:

```
1 xe host-evacuate
```

Evacuating all VMs from a host leaves the distribution of migrated VMs to XenServer.

- e) Shut down the host.

```
1 xe host-shutdown
```

- f) Boot the host using the XenServer installation media and method of your choice (such as, USB or network).
- g) Follow the XenServer installation procedure until the installer offers you the option to upgrade. Choose to upgrade.
- h) After the host upgrade is complete, start or resume any shutdown or suspended VMs.
- i) Migrate any VMs that you want back to the host.

If the upgrade of a subordinate host fails or is interrupted, you do not have to revert. Run the command `xe host-forget` in the pool to forget that host. Reinstall XenServer on the host, and then join it, as a new host, to the pool using the command `xe pool-join`.

- 8. After the XenServer hosts are updated, take the machines out of maintenance mode in Citrix Studio or Web Studio.

Mixed-mode pools

A mixed-mode pool is one where hosts in the pool are using different versions of XenServer. Do not operate your pool in mixed-mode (with multiple versions of XenServer) for longer than necessary, as the pool operates in a degraded state during upgrade. In this degraded state, certain VM, SR, VDI, and host operations are blocked. VMs that have run on a host at the higher version of XenServer cannot be migrated to or started on a host at the lower version of XenServer.

Mixed-mode pools are not supported for standard usage and are only supported as a transitional state during the upgrade of a pool. If you experience an issue while running in mixed mode, Technical Support will ask you to complete your pool upgrade and then reproduce the issue in a non-mixed pool.

After reviewing the upgrade options for your Citrix Virtual Apps and Desktops environment, your planned XenServer upgrade path might take longer than the available maintenance window. If possible, extend the maintenance window to enable your XenServer upgrade to complete within it. If this is not possible, you can choose to run the pool in mixed mode until your next maintenance window. However, running your pool in mixed mode increases the likelihood of unexpected behaviors or issues that might cause you to need an emergency maintenance window instead. Plan to minimize the time your pool spends in mixed mode.

If your Citrix Virtual Apps and Desktops environment is running temporarily on top of a mixed-mode XenServer pool, be aware of the following behaviors:

- For Pooled Desktop workloads that require the VMs to restart before they are reused, the VMs are restarted only on the hosts that are running the newer version of XenServer. The effective capacity of the pool is restricted. Depending on how many of the hosts in your pool have been upgraded, there might be insufficient capacity for all required VMs to be restarted. This behavior can result in failures and some Citrix Virtual Apps and Desktops users might not be able to access their required sessions.
- If you have dedicated machines using local storage that are located on hosts running the older version of XenServer, these VMs can be stopped, but they cannot be restarted until the upgrade is complete and the pool is no longer in mixed mode.

IntelliCache™

December 17, 2025

Using XenServer® with *IntelliCache* makes hosted Virtual Desktop Infrastructure deployments more cost-effective by enabling you to use a combination of shared storage and local storage. It is of particular benefit when many Virtual Machines (VMs) all share a common OS image (the primary image). The load on the storage array is reduced and performance is enhanced. In addition, network traffic to and from shared storage is reduced as the local storage caches the primary image from shared storage.

IntelliCache works by caching data for a VM in local storage on the VM host. This local cached data includes:

- A cached copy of the primary image (a VM's parent VDI). This file is thin provisioned and read only. It is populated as data is read from the primary image VDI on shared storage. When many VMs share a common primary image VDI, a VM can use the data read into the cache from another VM. Further access to the primary image on shared storage is not required.

- A cached copy of the per-VDI delta disk. Changes can be written to this local delta disk. For non-persistent VMs, data in the local cached delta disk is discarded on VM boot. For persistent VMs, data in the local cached delta disk is written back to the delta disk on the shared storage on VM boot.

[This is an embedded video. Click the link to watch the video](#)

Prerequisites

The IntelliCache feature has the following requirements:

- **Premium Edition license:** IntelliCache is only supported when using XenServer Premium Edition with Citrix Virtual Desktops.
- **Storage:** To use IntelliCache, you require both a local SR for the cached data and a shared SR for the source Virtual Disk Image (VDI). Use these SRs only for IntelliCache VMs. Do not mix traditional virtual machines and IntelliCache VMs on the same SR.
 - **Shared storage:** The shared storage used to host the source VDI must be one of the following types:
 - NFS
 - SMB3
 - Shared LVM (iSCSI or HBA)
 - **Local storage:** We recommend that you use a high performance local storage device to ensure the fastest possible data transfer. For example, use a Solid State Disk or a high performance RAID array. Consider both data throughput and storage capacity when sizing local disks.

A thin-provisioned, EXT3/EXT4 local SR is required for IntelliCache. Thin provisioning **must be** enabled for Citrix Virtual Desktops local caching to work properly. You can enable thin provisioning at XenServer install time or later, at the command line. Thin provisioning changes the default local storage type of the host from LVM to EXT4. For more information, see IntelliCache deployment.

Thin provisioning is a way of optimizing the use of available storage. This approach allows you to make more use of local storage instead of shared storage. It relies on on-demand allocation of blocks of data. In other approaches, all blocks are allocated up front. Thin provisioning allows the administrator to present more storage space to the VMs connecting to the Storage Repository (SR) than is available on the SR. There are no space guarantees, and allocation of a LUN does not claim any data blocks until the VM writes data.

Thin-provisioned SRs might run out of physical space, as the VMs within can grow to consume disk capacity on demand. IntelliCache VMs handle this condition by automatically

falling back to shared storage when the local SR cache is full. Do not mix traditional virtual machines and IntelliCache VMs on the same SR, as IntelliCache VMs can grow quickly in size.

IntelliCache deployment

To use the Intellicache feature, you must enable it both on XenServer and on Citrix Virtual Apps and Desktops™.

On XenServer

To use IntelliCache, thin provisioning of the local SR must be enabled either during host installation or be enabled manually on a running host using the CLI. Thin Provisioning changes the default local storage type of the host from LVM to EXT4.

Enable on host installation To enable thin provisioning during host installation, on the **Virtual Machine Storage** screen, select **Enable thin provisioning**. This option selects the host's local SR to be the one to be used for the local caching of VM VDIs.



Convert an existing host to use thin provisioning To delete an existing LVM local SR, and replace it with a thin-provisioned EXT3/EXT4 SR, enter the following commands.

Warning:

These commands remove your existing local SR, and VMs on the SR are permanently deleted.


```
1    localsr=`xe sr-list type=lvm host=hostname params=uuid --minimal`
2    echo localsr=$localsr
3    pbd=`xe pbd-list sr-uuid=$localsr params=uuid --minimal`
4    echo pbd=$pbd
5    xe pbd-unplug uuid=$pbd
6    xe pbd-destroy uuid=$pbd
7    xe sr-forget uuid=$localsr
8    sed -i "s/'lvm'/'ext'/" /etc/firstboot.d/data/default-storage.
    conf
9    rm -f /var/lib/misc/ran-storage-init
10   systemctl restart storage-init.service
11   xe sr-list type=ext
```

To enable local caching, enter the following commands:

```
1    xe host-disable host=hostname
2    localsr=`xe sr-list type=ext host=hostname params=uuid --
    minimal`
3    xe host-enable-local-storage-caching host=hostname sr-uuid=
    $localsr
4    xe host-enable host=hostname
```

On Citrix Virtual Apps and Desktops

When connecting from Citrix Virtual Apps and Desktops to XenServer in Citrix Studio, select to use IntelliCache. If thin provisioning of the local SR has been enabled in XenServer, Citrix Studio shows the option **Use IntelliCache to reduce load on the shared storage device** when creating the connection to XenServer.

For more information, see the [Citrix Virtual Apps and Desktops product documentation](#).

VM behavior with IntelliCache

The VDI flag `on-boot` dictates the behavior of a VM VDI when the VM is booted and the VDI flag `allow-caching` dictates the caching behavior.

The values to use for these parameters depends on the type of VM you are creating and what its intended use is:

- **For shared or randomly allocated machines:**
 - Set the `on-boot` parameter to `reset`.
 - Set the `allow-caching` parameter to `true`

For example:

```
1    xe vdi-param-set uuid=vdi_uuid on-boot=reset allow-caching=true
```

On VM boot, the VDI is reverted to the state it was in at the previous boot. All changes while the VM is running are lost when the VM is next booted. New VM data is written only to local storage. There are no writes to shared storage. This approach means that the load on shared storage is reduced. However the VM cannot be migrated between hosts.

Select this option if you plan to deliver standardized desktops to which users cannot make permanent changes.

- **For static or dedicated machines:**

- Set the `on-boot` parameter to `persist`.
- Set the `allow-caching` parameter to `true`

For example:

```
1 xe vdi-param-set uuid=vdi_uuid on-boot=persist allow-caching=true
```

On VM boot, the VDI is in the state it was left in at the last shutdown. New VM data is written to both local and shared storage. Reads of cached data do not require I/O traffic to shared storage so the load on shared storage is reduced. VM migration to another host is permitted and the local cache on the new host is populated as data is read.

Select this option if you plan to allow users to make permanent changes to their desktops.

Note:

For VMs whose VDIs are on a GFS2 SR, the VM on-boot behavior is different to VMs with VDIs on other types of SRs. For VDIs on a GFS2 SR, the on-boot option is applied on VM shutdown, not on VM boot.

Implementation details and troubleshooting

Q: Is IntelliCache compatible with live migration and high availability?

A: You can use live migration and high availability with IntelliCache when virtual desktops are in Private mode, that is, when `on-boot=persist`.

Warning:

A VM cannot be migrated if any of its VDIs have caching behavior flags set to `on-boot=reset` and `allow-caching=true`. Migration attempts for VMs with these properties fail.

Q: Where does the local cache live on the local disk?

A: The cache lives in a Storage Repository (SR). Each host has a configuration parameter (called `local-cache-sr`) indicating which (local) SR is to be used for the cache files. Typically, this SR is an EXT3/EXT4

type SR. When you run VMs with IntelliCache, you see files inside the SR with names `uuid.vhdcache`. This file is the cache file for the VDI with the given UUID. These files are not displayed in XenCenter. The only way of seeing them is by logging into dom0 and listing the contents of `/var/run/sr-mount/sr-uuid`

Q: How do I specify a particular SR for use as the cache?

A: The host object field `local-cache-sr` references a local SR. You can view its value by running the following command:

```
1 xe sr-list params=local-cache-sr,uuid,name-label
```

This field is set either:

- After host installation, if you have chosen the “Enable thin provisioning” option in the host installer, or
- By running `xe host-enable-local-storage-caching host=host sr-uuid=sr`. The command requires the specified host to be disabled. Shut down the VMs when you use this command.

The first option uses the EXT3/EXT4 type local SR and is created during host installation. The second option uses the EXT3/EXT4 type SR that is specified on the command line.

Warning:

These steps are only necessary for users who have configured more than one local SR.

Q: When is the local cache deleted?

A: The local cache includes two types of information:

- A shared cache for the primary image disk. This shared cache file mirrors the read-only VDI from the shared storage and is not deleted or reset when actions are taken on individual VDIs.
- A per-VDI delta disk cache for each VDI. For a non-persistent VM, the delta disk cache is reset when a VDI is attached to a VM (for example, on VM start). This cache file is only deleted when the VDI itself is deleted. If the host is offline when you delete the VDI, the SR synchronization that runs on startup garbage collects the cache file.

Note:

The delta disk cache file is not deleted from the host when a VM migrates to a different host or is shut down.

PVS-Accelerator

February 18, 2026

The XenServer PVS-Accelerator feature offers extended capabilities for customers using XenServer with Citrix Provisioning. Citrix Provisioning is a popular choice for image management and hosting for Citrix Virtual Apps and Desktops™ or Citrix DaaS. PVS-Accelerator dramatically improves the already excellent combination of XenServer and Citrix Provisioning. Some of the benefits that this new feature provides include:

- **Data locality:** Use the performance and locality of memory, SSD, and NVM devices for read requests, while substantially reducing network utilization.
- **Improved end-user experience:** Data locality enables a reduction in the read I/O latency for cached target devices (VMs), further accelerating end-user applications.
- **Accelerated VM boots and boot storms:** Reduced read I/O-latency and improved efficiency can accelerate VM boot times and enable faster performance when many devices boot up within a narrow time frame.
- **Simplified scale-out by adding more hypervisor hosts:** Fewer Citrix Provisioning™ servers may be needed as the storage load is efficiently dispersed across all XenServer hosts. Peak loads are handled using the cache within originating hosts.
- **Reduced TCO and simplified infrastructure requirements:** Fewer Citrix Provisioning servers means a reduction in hardware and license requirements, in addition to reduced management overhead. Freed up capacity is available for workloads.

Notes:

PVS-Accelerator is available for XenServer Premium Edition customers. To use the PVS-Accelerator feature, upgrade the Citrix License Server to version 11.14 or later.

To use PVS-Accelerator with UEFI-enabled VMs, ensure that you are using Citrix Provisioning 1906 or later.

[This is an embedded video. Click the link to watch the video](#)

How does PVS-Accelerator work

PVS-Accelerator employs a Proxy mechanism that resides in the Control Domain (dom0) of XenServer. When this feature is enabled, Citrix Provisioning targets device (VM) read requests are cached directly on the XenServer host machine. These requests are cached in physical memory or a storage repository. When subsequent VMs on that XenServer host make the same read request, the virtual disk is

streamed directly from cache, not from the Citrix Provisioning server. Removing the need to stream from the Citrix Provisioning server reduces network utilization and processing on the server considerably. This approach results in a substantial improvement in VM performance.

Considerations

Consider the following when using the PVS-Accelerator feature:

- Citrix Provisioning target devices are aware of their proxy status. No additional configuration is required once the capability is installed.
- The PVS-Accelerator supports both IPv4 and IPv6-based protocols. (IPv6 requires Citrix Provisioning 2402 CU1, Citrix Provisioning 2407, or later.)
- In environments where multiple Citrix Provisioning servers are deployed with the same VHD, but have different file system timestamps, data might be cached multiple times. Due to this limitation, we recommend using VHDX format, rather than VHD for virtual disks.
- Do not use a large port range for PVS server communication. Setting a range of more than 20 ports on the PVS server is rarely necessary. A large port range can slow packet processing and increase the boot time of VMs when using PVS-Accelerator.
- After you start a VM with PVS-Accelerator enabled, the caching status for the VM is displayed in XenCenter®:
 - In the **PVS** tab of the pool or the host
 - In the **General** tab for the VM
- You cannot run more than 200 PVS-Accelerator-enabled VMs on a XenServer host.
- Customers can confirm the correct operation of the PVS-Accelerator using RRD metrics on the host's **Performance** tab in XenCenter. For more information, see [Monitor and manage your deployment](#).
- PVS-Accelerator requires Citrix Provisioning 7.13 or later.
- To use PVS-Accelerator with UEFI-enabled VMs, ensure that you are using Citrix Provisioning 1906 or later.
- PVS-Accelerator is available for XenServer Premium Edition customers.
- PVS-Accelerator requires License Server 11.14 or later.
- PVS-Accelerator works on the first virtual network interface (VIF) of a cached VM. Therefore, connect the first VIF to the Citrix Provisioning storage network for caching to work.
- PVS-Accelerator can currently not be used on network ports which enforce that IPs are bound to certain MAC addresses. This switch functionality might be called "IP Source Guard" or similar.

In such environments, PVS targets fail to boot with error ‘Login request time out!’ after enabling PVS-Accelerator.

Enable PVS-Accelerator

Customers must complete the following configuration settings in XenServer and in Citrix Provisioning to enable the PVS-Accelerator feature:

1. Configure PVS-Accelerator in XenServer by using XenCenter or the xe CLI. This configuration includes adding a Citrix Provisioning site and specifying the location for Citrix Provisioning cache storage.
 - For CLI instructions, see *Configuring PVS-Accelerator in XenServer by using the CLI* in the following section.
 - For information about configuring PVS-Accelerator using XenCenter, see [PVS-Accelerator](#) in the XenCenter documentation.
2. After configuring PVS-Accelerator in XenServer, complete the cache configuration for the PVS Site using the PVS UI. For detailed instructions, see [Completing the cache configuration in Citrix Provisioning](#).

Configuring ports

Citrix Provisioning Services uses the following ports:

- 6901, 6902, 6905: Used for provisioning server outbound communication (packets destined for the target device)
- 6910: Used for target device logon with Citrix Provisioning Services
- Configurable target device port. The default port is 6901.
- Configurable server port range. The default range is 6910-6930.

For information about the ports used by Citrix Provisioning Services, see [Communication ports used by XenServer](#).

The configured port range in XenServer must include all the ports in use. For example, use 6901-6930 for the default configuration.

Note:

Do not use a large port range for PVS server communication. Setting a range of more than 20 ports on the PVS server is rarely necessary. A large port range can slow packet processing and increase the boot time of VMs when using PVS-Accelerator.

Configure PVS-Accelerator in XenServer by using the CLI

1. Run the following command to create a Citrix Provisioning site configuration on XenServer:

```
1 PVS_SITE_UUID=$(xe pvs-site-introduce name=label=My PVS Site)
```

2. For each host in the pool, specify what cache to use. You can choose to store the cache on a storage repository (SR) or in the Control Domain Memory.

Configure cache storage on a storage repository Consider the following characteristics when choosing a storage repository (SR) for cache storage:

Advantages:

- Most recently read data is cached in the memory on a best effort basis. Accessing the data can be as fast as using the Control Domain memory.
- The cache can be much larger when it is on an SR. The cost of the SR space is typically a fraction of the cost of the memory space. Caching on an SR can take more load off the Citrix Provisioning server.
- You don't have to modify the Control Domain memory setting. The cache automatically uses the memory available in the Control Domain and never causes the Control Domain to run out of memory.
- The cache VDI's can be stored on shared storage. However, this choice of storage rarely makes sense. This approach only makes sense where the shared storage is significantly faster than the Citrix Provisioning server.
- You can use either a file-based or a block-based SR for cache storage.

Disadvantages:

- If the SR is slow and the requested data isn't in the memory tier, the caching process can be slower than a remote Citrix Provisioning server.
- Cached VDI's that are stored on shared storage cannot be shared between hosts. A cached VDI is specific to one host.

Perform the following steps to configure cache storage on a Storage Repository:

1. Run the following command to find the UUID of the SR that to use for caching:

```
1 xe sr-list name=label=Local storage host=host-name-label --minimal  
  )
```

2. Create the cache-storage.

```
1 xe pvs-cache-storage-create host=host-name-label pvs-site-uuid=  
  PVS_SITE_UUID \  
2      sr-uuid=SR_UUID size=10GiB
```

Note:

When selecting a Storage Repository (SR), the feature uses up to the specified cache size on the SR. It also implicitly uses available Control Domain memory as a best effort cache tier.

Configuring cache storage in the control domain memory Consider the following characteristics when choosing the Control Domain memory for cache storage:

Advantages:

Using memory means consistently fast Read/Write performance when accessing or populating the cache.

Disadvantages:

- Hardware must be sized appropriately as the RAM used for cache storage is not available for VMs.
- Control Domain memory must be extended **before** configuring cache storage.

Note:

If you choose to store the cache in the Control Domain memory, the feature uses up to the specified cache size in Control Domain memory. This option is only available after extra memory has been assigned to the Control Domain. For information about increasing the Control Domain memory, see [Change the amount of memory allocated to the control domain](#).

After you increase the amount of memory allocated to the Control Domain of the host, the additional memory can be explicitly assigned for PVS-Accelerator.

Perform the following steps to configure cache storage in the Control Domain memory:

1. Run the following command to find the UUID of the host to configure for caching:

```
1 xe host-list name=label=host-name-label --minimal
```

2. Create an SR of the special type `tmpfs`:

```
1 xe sr-create type=tmpfs name=label=MemorySR host-uuid=  
HOST_UUID device-config:uri=""
```

Note:

For SRs of the special type `tmpfs`, the value of the required parameter `name=label` is disregarded and a fixed name is used instead.

3. Run the following command to create the cache storage:

```
1 xe pvs-cache-storage-create host-uuid=HOST_UUID  
2 pvs-site-uuid=PVS_SITE_UUID sr-uuid=SR_UUID size=1GiB
```

Where **SR_UUID** is the UUID of the SR created in step b

Complete the cache configuration in Citrix Provisioning

After configuring PVS-Accelerator in XenServer, perform the following steps to complete the cache configuration for the Citrix Provisioning site.

In the Citrix Provisioning Administrator Console, use the Citrix Virtual Desktops™ Setup Wizard or the Streaming VM Wizard (depending on your deployment type) to access the Proxy capability. Although both wizards are similar and share many of the same screens, the following differences exist:

- The **Citrix Virtual Desktops Setup Wizard** is used to configure VMs running on XenServer hypervisor that is controlled using Citrix Virtual Desktops.
- The **Streaming VM Wizard** is used to create VMs on a host. It does not involve Citrix Virtual Desktops.

Launch the Citrix Provisioning Administrator Console:

1. Navigate to the Citrix Provisioning site.
2. Select the Citrix Provisioning site, right-click to expose a contextual menu.
3. Choose the appropriate wizard based on the deployment. Select the option **Enable PVS-Accelerator for all Virtual Machines** to enable the PVS-Accelerator feature.
4. If you are enabling virtual disk caching for the first time, the **XenServer** screen appears on the Streamed Virtual Machine Setup wizard. It displays the list of all Citrix Provisioning sites configured on XenServer that have not yet been associated with a Citrix Provisioning site. Using the list, select a Citrix Provisioning site to apply PVS-Accelerator. This screen is not displayed when you run the wizard for the same Citrix Provisioning site using the same XenServer host.
5. Click **Next** to complete the caching configuration.
6. Click **Finish** to provision Citrix Virtual Desktops or Streamed VMs and associate the selected Citrix Provisioning site with the PVS Accelerator in XenServer. When this step is complete, the **View PVS Servers** button in the **PVS-Accelerator configuration** window is enabled in XenCenter. Clicking the **View PVS Servers** button displays the IP addresses of all PVS Servers associated with the Citrix Provisioning site.

Caching operation

The PVS-Accelerator functionality caches:

- **Reads** from virtual disks but not writes or reads from a write cache
- **Based on image versions.** Multiple VMs share cached blocks when they use the same image version
- Devices with any **non-persistent** write cache type
- Virtual disks with the **access mode Standard Image**. It does not work for virtual disks with the access mode Private Image
- Devices that are marked as **type Production or Test**. Devices marked as type Maintenance are not cached

PVS-Accelerator CLI operations

The following section describes the operations that customers can perform when using PVS-Accelerator using the CLI. Customers can also perform these operations using XenCenter. For more information, see [PVS-Accelerator](#) in the XenCenter documentation.

View Citrix Provisioning server addresses and ports configured by Citrix Provisioning

PVS-Accelerator works by optimizing the network traffic between a VM and the Citrix Provisioning server. When completing the configuration on the Citrix Provisioning server, the Citrix Provisioning server populates the `pvs-server` objects on XenServer with their IPs and ports. PVS-Accelerator later uses this information to optimize specifically the traffic between a VM and its Citrix Provisioning servers. The configured Citrix Provisioning servers can be listed using the following command:

```
1 xe pvs-server-list pvs-site-uuid=PVS_SITE_UUID params=all
```

Configure a VM for caching

PVS-Accelerator can be enabled for the VM by using any of the following tools:

- Citrix Provisioning CLI
- Citrix Virtual Desktops Setup Wizard
- Streamed VM Setup Wizard
- XenCenter
- The xe CLI

The xe CLI configures PVS-Accelerator by using the VIF of a VM. It creates a Citrix Provisioning proxy that links the VM's VIF with a Citrix Provisioning site.

To configure a VM:

1. Find the first VIF of the VM to enable caching on it:

```
1 VIF_UUID=$(xe vif-list vm-name=label=pvsdevice_1 device=0 --minimal)
```

2. Create the Citrix Provisioning proxy

```
1 xe pvs-proxy-create pvs-site-uuid=PVS_SITE_UUID vif-uuid=$VIF_UUID
```

Disable caching for a VM

PVS-Accelerator can be disabled for a VM by destroying the Citrix Provisioning proxy that links the VM's VIF with a `pvs-site`.

1. Find the first VIF of the VM:

```
1 VIF_UUID=$(xe vif-list vm-name=label=pvsdevice_1 device=0 --minimal)
```

2. Find the Citrix Provisioning proxy of the VM:

```
1 PVS_PROXY_UUID=$(xe pvs-proxy-list vif-uuid=$VIF_UUID --minimal)
```

3. Destroy the Citrix Provisioning proxy:

```
1 xe pvs-proxy-destroy uuid=$PVS_PROXY_UUID
```

Remove the PVS-Accelerator storage for a host or a site

To remove the PVS-Accelerator storage for a host or a site:

1. Find the host for which you would like to destroy the storage:

```
1 HOST_UUID=$(xe host-list name=label=HOST_NAME --minimal)
```

2. Find the UUID of the object:

```
1 PVS_CACHE_STORAGE_UUID=$(xe pvs-cache-storage-list host-uuid=$HOST_UUID --minimal)
```

3. Destroy the object:

```
1 xe pvs-cache-storage-destroy uuid=$PVS_CACHE_STORAGE_UUID
```

Forget the PVS-Accelerator configuration for a site

To forget the PVS-Accelerator configuration for a site:

1. Find the Citrix Provisioning site:

```
1 PVS_SITE_UUID=$(xe pvs-site-list name-label=My PVS Site)
```

2. Run the following command to forget the Citrix Provisioning site:

```
1 xe pvs-site-forget uuid=$PVS_SITE_UUID
```

Resource pools

June 30, 2026

A *resource pool* comprises multiple XenServer® host installations, bound together to a single managed entity that can host virtual machines. If combined with shared storage, a resource pool enables VMs to be started on *any* XenServer host which has sufficient memory.

The *pool coordinator* (formerly “pool master”) is a server in the resource pool that exposes an administration interface (used by XenCenter® and the XenServer command line interface, known as the `xe CLI`). The pool coordinator forwards commands to individual members as necessary.

This article describes the concepts, requirements, and best practices with regard to resource pools. For information about how to create and manage your pools, see [Manage your pools](#).

Advantages of resource pools

While you can organize your XenServer hosts as *standalone* hosts, effectively a pool of one, XenServer is optimized for hosts grouped into resource pools with shared storage. Several of our features, such as high availability and live migration, are only available in resource pools of multiple hosts. The advantages of organizing your XenServer hosts into resource pools include the following:

- **VM mobility:** When your VMs are running on a pool and have their disks on the pool shared storage, these VMs can be dynamically moved among XenServer hosts while still running ([live migration](#)). Also, if an individual XenServer host suffers a hardware failure, the administrator can restart failed VMs on another XenServer host in the same resource pool.
- **High availability** This feature works to protect the VMs in your workload and ensure that they have minimal downtime in cases of hardware or host failure. When high availability is enabled on the resource pool, VMs can automatically restart on another host when their host fails. If the

pool coordinator fails, high availability elects another pool coordinator. For more information, see [High availability](#).

- **Workload balancing** Workload Balancing can evaluate your resource pool and the VM workload running on it to recommend the optimum placement for your VMs. You can also choose to have Workload Balancing automatically follow its placement recommendations and migrate your VMs between hosts in the pool. For more information, see [Workload Balancing](#).
- **Anti-affinity VM placement:** Ensure that the VMs in a group are evenly spread across the hosts in a pool. For more information, see [VM placement](#).
- **Ease of management:** Rather than managing each XenServer host individually, add the resource pool to XenCenter and manage all the hosts, SRs, and VMs within it together. For more information, see [XenCenter](#).

Requirements for creating resource pools

When designing your XenServer deployment and deciding about your pool configuration, consider the following requirements:

Hardware requirements

All of the servers in a XenServer resource pool must have broadly compatible CPUs, that is:

- The CPU vendor (Intel, AMD) must be the same on all CPUs on all servers.
- All CPUs must have virtualization enabled.

Depending on how similar the CPUs are, the pool is on one of the following types:

- **Homogeneous pool:** A homogeneous resource pool is an aggregate of servers with identical CPUs. CPUs on a server joining a homogeneous resource pool must have the same vendor, model, and features as the CPUs on servers already in the pool.
- **Heterogeneous pool:** Heterogeneous pool creation is made possible by using technologies in Intel (FlexMigration) and AMD (Extended Migration) CPUs that provide CPU *masking* or *leveling*. These features allow a CPU to be configured to *appear* as providing a different make, model, or feature set than it actually does. These capabilities enable you to create pools of hosts with different CPUs but still safely support live migrations. As a result of this feature masking or leveling, you might not get the full performance of your CPUs.

Joining host requirements

XenServer checks that the following conditions are true for the host joining the pool:

- It must be running the same version of XenServer, at the same update level, as hosts already in the pool.
- The joining host is not a member of an existing resource pool.
- The joining host has no shared storage configured.
- The joining host is not hosting any running or suspended VMs.
- No active operations are in progress on the VMs on the joining host, such as a VM shutting down or being exported.
- The clock on the joining host is synchronized to the same time as the pool coordinator (for example, by using NTP).
- The management interface of the joining host is not bonded. You can configure the management interface when the host successfully joins the pool.
- The joining host's management IP address is static, either configured on the host itself or by using an appropriate configuration on your DHCP server.
- The joining host's management interface is on the same tagged VLAN as that of the resource pool.
- The joining host must be configured with the same supplemental packs at the same revision as the hosts already in the pool.
- The joining host must have the same XenServer license as the hosts already in the pool. You can change the license of any pool members after joining the pool. The host with the lowest license determines the features available to all members in the pool.

If you add a host to a pool by using XenCenter, XenCenter ensures that a matching license is applied to the joining host. However, if you are doing a pool join by using the xe CLI, we recommend that you assign a matching license to the host before joining it to the pool.

- The joining host must be in the same site as the pool and connected by a low latency network.

Storage requirements

Shared storage used by the resource pool has the following requirements:

- Servers providing shared NFS or iSCSI storage for the pool must have a static IP address or a static DHCP lease.

Recommended pool size

The listed configuration limit of 32 is the maximum number of hosts that we support in a pool. However, it is not a recommended pool size, as it is often not the optimal size from a management or

performance point of view for most workloads. Consider the following factors when deciding on the best size for your deployment:

- **Management considerations:** Most XenServer management is performed at the pool level. A larger pool reduces the amount of management required and enables you to manage more hosts together. However, some management operations, for example applying updates to a pool, can take longer if you have more hosts, as the operation must be performed sequentially on each host in the pool. In this case, you might need a longer maintenance window to complete certain operations in a large pool than you need for multiple smaller pools.
- **Resource sharing:** A XenServer pool typically shares resources such as storage repositories between hosts. A larger pool allows more hosts to share resources, which can have benefits. For example, in your Citrix Virtual Apps and Desktops™ environment you can have more hosts sharing a gold image, rather than having to create multiple copies on different pools. However, the particular devices that you choose for your shared resources might have specific performance considerations that make it better to use smaller groupings of hosts.
- **Control plane performance:** In a XenServer pool, all operations are managed by the pool coordinator. The load on the toolstack of this host increases as you add more hosts to the pool: there are more background activities from each host and the expected number of concurrent operations increases. As the load on the toolstack increases, the time taken for each operation is likely to increase. As a result, one large pool can perform noticeably slower than two smaller pools.
- **Fault isolation:** If XenServer or another component that is critical to the pool (for example, a storage device) has a problem, in a larger pool this might have a greater impact on your workloads than if that workload is split over multiple smaller pools.
- **GFS2 storage:** If you use GFS2 storage for thin provisioning on block storage, XenServer supports a maximum of 16 hosts. This is due to limitations in the GFS2 implementation and the increased communications required between hosts to manage the storage.
- **High availability:** If you use our high availability functionality to protect your VMs, all hosts in the pool continually monitor each other and communicate their status. As the pool increases in size, the volume of messages that each host has to send and receive as part of this monitoring increases. If there is a high load on the control domain, this can increase the probability of some of these monitoring messages being lost. In extreme scenarios, lost messages can cause hosts to unexpectedly fence for safety. When using the high availability feature, we recommend that you use smaller pools, with a maximum of 16 hosts, to reduce the risk of unexpected fencing.

For most Citrix Virtual Apps and Desktops use cases, we recommend a pool size of 16 hosts, with a maximum size of 32.

In addition to considering these factors when planning your pool size, observe and monitor your pool behavior to determine whether you need to change the pool size in your running environment.

Communicate with XenServer hosts and resource pools

TLS

XenServer uses the TLS 1.2 protocol to encrypt management API traffic. Any communication between XenServer and management API clients (or appliances) uses the TLS 1.2 protocol.

Important:

We do not support customer modifications to the cryptographic functionality of the product.

XenServer uses the following cipher suites:

- ECDHE-RSA-AES256-GCM-SHA384
- ECDHE-RSA-AES128-GCM-SHA256

SSH

When using an SSH client to connect directly to the XenServer host the following algorithms can be used:

Ciphers:

- aes128-ctr
- aes256-ctr
- aes128-gcm@openssh.com
- aes256-gcm@openssh.com

MACs:

- hmac-sha2-256
- hmac-sha2-512
- hmac-sha1

KexAlgorithms:

- curve25519-sha256
- ecdh-sha2-nistp256
- ecdh-sha2-nistp384
- ecdh-sha2-nistp521

HostKeyAlgorithms:

- ecdsa-sha2-nistp256
- ecdsa-sha2-nistp384

- ecdsa-sha2-nistp521
- ssh-ed25519
- ssh-rsa

Important:

We do not support customer modifications to the cryptographic functionality of the product. However, if you want to disable SSH access to a XenServer host, see [Disable SSH access for a host](#) or [Disable SSH access for a pool](#).

What happens when a host joins a resource pool?

When a new host joins a resource pool, the joining host synchronizes its local database with the pool-wide one, and inherits some settings from the pool:

- VM, local, and remote storage configuration are added to the pool-wide database. This configuration is applied to the joining host in the pool unless you explicitly make the resources shared after the host joins the pool.
- The joining host inherits existing shared storage repositories in the pool. Appropriate PBD records are created so that the new host can access existing shared storage automatically.
- Networking information is partially inherited to the joining host: the *structural* details of NICs, VLANs, and bonded interfaces are all inherited, but *policy* information is not. This policy information, which must be reconfigured, includes:
 - The IP addresses of the management NICs, which are preserved from the original configuration.
 - The location of the management interface, which remains the same as the original configuration. For example, if the other pool hosts have management interfaces on a bonded interface, the joining host must be migrated to the bond after joining.
 - Dedicated storage NICs, which must be reassigned to the joining host from XenCenter or the CLI, and the PBDs replugged to route the traffic accordingly. This is because IP addresses are not assigned as part of the pool join operation, and the storage NIC works only when this is correctly configured. For more information on how to dedicate a storage NIC from the CLI, see [Manage networking](#).

Manage your pools

June 10, 2026

This article describes some of the actions you can take to manage your pool.

To manage individual hosts, see [Manage your hosts](#).

Create a resource pool

Resource pools can be created using XenCenter® or the CLI. When a new host joins a resource pool, the joining host synchronizes its local database with the pool-wide one, and inherits some settings from the pool:

- VM, local, and remote storage configuration are added to the pool-wide database. This configuration is applied to the joining host in the pool unless you explicitly make the resources shared after the host joins the pool.
- The joining host inherits existing shared storage repositories in the pool. Appropriate PBD records are created so that the new host can access existing shared storage automatically.
- Networking information is partially inherited to the joining host: the *structural* details of NICs, VLANs, and bonded interfaces are all inherited, but *policy* information is not. This policy information, which must be reconfigured, includes:
 - The IP addresses of the management NICs, which are preserved from the original configuration.
 - The location of the management interface, which remains the same as the original configuration. For example, if the other pool hosts have management interfaces on a bonded interface, the joining host must be migrated to the bond after joining.
 - Dedicated storage NICs, which must be reassigned to the joining host from XenCenter or the CLI, and the PBDs replugged to route the traffic accordingly. This is because IP addresses are not assigned as part of the pool join operation, and the storage NIC works only when this is correctly configured. For more information on how to dedicate a storage NIC from the CLI, see [Manage networking](#).
 - Reconfigure the management interface and move it on to a physical NIC before adding the host to the pool. After the host has joined the pool, you can reconfigure the management interface again.

Before creating a resource pool, review the requirements for the pool and the joining host. For more information, see [Resource pools](#).

Add a host to a pool by using the xe CLI

1. Update your pool and the joining host to the same level before attempting the join. You can do this in one of the following ways:

- Update the joining host to the same level as the pool:
 - (Recommended) By synchronizing the joining host with the pool and updating from update files held on the pool coordinator. For more information, see [Update to pool level](#).
 - By downloading an update bundle that matches the pool level and using it to update the joining host.
- Update both the pool and the joining host to the latest level:
 - By synchronizing the host and pool with the same update channel and applying the latest updates.
 - By using the offline updates mechanism to apply the latest level of updates to both the host and pool.

For more information about these update methods see [Apply updates by using XenCenter](#) or [Apply updates by using the xe CLI](#).

2. Open a console on the XenServer® host that you want to join to a pool.
3. Join the XenServer host to the pool by issuing the command:

```
1 xe pool-join master-address=<address of pool coordinator> master-username=<administrator username> master-password=<password>
```

The `master-address` must be set to the fully qualified domain name of the pool coordinator. The `password` must be the administrator password set when the pool coordinator was installed.

Note:

When you join a host to a pool, the administrator password for the joining host is automatically changed to match the administrator password of the pool coordinator.

XenServer hosts belong to an unnamed pool by default. To create your first resource pool, rename the existing nameless pool. Use tab-complete to find the `pool_uuid`:

```
1 xe pool-param-set name-label="New Pool" uuid=pool_uuid
```

Create heterogeneous resource pools

XenServer simplifies expanding deployments over time by allowing disparate host hardware to be joined in to a resource pool, known as heterogeneous resource pools. Heterogeneous resource pools are made possible by using technologies in Intel (FlexMigration) and AMD (Extended Migration) CPUs that provide CPU “masking” or “leveling”. The CPU masking and leveling features allow a CPU to be configured to *appear* as providing a different make, model, or functionality than it actually does. This

feature enables you to create pools of hosts with disparate CPUs but still safely support live migration.

Note:

The CPUs of XenServer hosts joining heterogeneous pools must be of the same vendor (that is, AMD, Intel) as the CPUs of the hosts already in the pool. However, the hosts are not required to be the same type at the level of family, model, or stepping numbers.

XenServer simplifies the support of heterogeneous pools. Hosts can now be added to existing resource pools, irrespective of the underlying CPU type (as long as the CPU is from the same vendor family). The pool feature set is dynamically calculated every time:

- A new host joins the pool
- A pool member leaves the pool
- A pool member reconnects following a reboot

Any change in the pool feature set does not affect VMs that are currently running in the pool. A Running VM continues to use the feature set which was applied when it was started. This feature set is fixed at boot and persists across migrate, suspend, and resume operations. If the pool level drops when a less-capable host joins the pool, a running VM can be migrated to any host in the pool, except the newly added host. When you move or migrate a VM to a different host within or across pools, XenServer compares the VM's feature set against the feature set of the destination host. If the feature sets are found to be compatible, the VM is allowed to migrate. This enables the VM to move freely within and across pools, regardless of the CPU features the VM is using. If you use Workload Balancing to select an optimal destination host to migrate your VM, a host with an incompatible feature set will not be recommended as the destination host.

Add shared storage

This section shows an example of how shared storage (represented as a storage repository) can be created on an existing NFS server. For a complete list of supported shared storage types, see [Storage](#).

To add NFS shared storage to a resource pool by using the CLI

1. Open a console on any XenServer host in the pool.
2. Create the storage repository on server:/path by issuing the following command:

```
1 xe sr-create content-type=user type=nfs name-label="Example SR"  
   shared=true \  
2   device-config:server=server \  
3   device-config:serverpath=path
```

`device-config:server` is the host name of the NFS server and `device-config:serverpath` is the path on the NFS server. As `shared` is set to true, shared storage is automatically connected to every XenServer host in the pool. Any XenServer hosts that join later are also connected to the storage. The Universally Unique Identifier (UUID) of the storage repository is printed on the screen.

3. Find the UUID of the pool by running the following command:

```
1 xe pool-list
```

4. Set the shared storage as the pool-wide default with the following command:

```
1 xe pool-param-set uuid=pool_uuid default-SR=sr_uuid
```

As the shared storage has been set as the pool-wide default, all future VMs have their disks created on shared storage by default. For information about creating other types of shared storage, see [Create an SR](#).

Remove XenServer hosts from a resource pool

Note:

Before removing any XenServer host from a pool, ensure that you shut down all the VMs running on that host. Otherwise, you can see a warning stating that the host cannot be removed.

When you remove (*eject*) a host from a pool, the machine is rebooted, reinitialized, and left in a state similar to a fresh installation. Do not eject XenServer hosts from a pool if there is important data on the local disks.

To remove a host from a resource pool by using the CLI

1. Open a console on any host in the pool.
2. Find the UUID of the host by running the following command:

```
1 xe host-list
```

3. Eject the required host from the pool:

```
1 xe pool-eject host-uuid=host_uuid
```

The XenServer host is ejected and left in a freshly installed state.

Warning:

Do *not* eject a host from a resource pool if it contains important data stored on its local

disks. All of the data is erased when a host is ejected from the pool. If you want to preserve this data, copy the VM to shared storage on the pool using XenCenter, or the `xe vm-copy` CLI command.

When XenServer hosts containing locally stored VMs are ejected from a pool, the VMs will be present in the pool database. The locally stored VMs are also visible to the other XenServer hosts. The VMs do not start until the virtual disks associated with them have been changed to point at shared storage seen by other XenServer hosts in the pool, or removed. Therefore, we recommend that you move any local storage to shared storage when joining a pool. Moving to shared storage allows individual XenServer hosts to be ejected (or physically fail) without loss of data.

Note:

When a host is removed from a pool that has its management interface on a tagged VLAN network, the machine is rebooted and its management interface will be available on the same network.

Rotate the pool secret

The pool secret is a secret shared among the hosts in a pool that enables the host to prove its membership to a pool.

Because users with the Pool Admin role can discover this secret, it is good practice to rotate the pool secret if one of these users leaves your organization or loses their Pool Admin role.

You can rotate the pool secret by using XenCenter or the `xe` CLI.

XenCenter

To rotate the pool secret for a pool by using XenCenter, complete the following steps:

1. In the **Resources** pane, select the pool or any host in the pool.
2. On the **Pool** menu, select **Rotate Pool Secret**.

When you rotate the pool secret, you are also prompted to change the root password. If you rotated the pool secret because you think that your environment has been compromised, ensure that you also change the root password of the pool coordinator. For more information, see [Change the password](#).

xe CLI

To rotate the pool secret by using the `xe` CLI, run the following command on a host in the pool:

```
1 xe pool-secret-rotate
```

If you rotated the pool secret because you think that your environment has been compromised, ensure that you also change the root password. For more information, see [Change the password](#).

Configure SSH access

SSH access to your XenServer pool is enabled by default. If you want to disable SSH access to your XenServer pool, run the following command:

```
1 xe pool-disable-ssh pool=<pool-uuid_or_name_label>
```

This command disables SSH. The pool refuses new SSH connections but does not disconnect any existing sessions.

To enable SSH access to your XenServer pool, run the following command:

```
1 xe pool-enable-ssh pool=<pool-uuid_or_name_label>
```

To set current pool SSH temporary enabled timeout (seconds):

```
1 xe pool-param-set uuid=<pool-uuid> ssh-enabled-timeout=<seconds>
```

The timeout parameter determines SSH service duration. When set to 0, SSH remains active indefinitely. When set to a positive integer (seconds), SSH is automatically disabled after the specified period. Configuration changes take immediate effect on following active SSH sessions and take effect every time the user enables SSH.

Note:

The maximum value you can set for `ssh-enabled-timeout` is 172800 seconds (which is approximately 2 days).

To set current pool's SSH console idle timeout (seconds)

```
1 xe pool-param-set uuid=<pool-uuid> console-idle-timeout=<seconds>
```

Configure the idle session timeout for VNC and SSH console connections using a non-negative integer value in seconds, where 0 disables timeout (sessions never expire) and positive values (for example, 3600) automatically terminate idle sessions after the specified duration, with settings applying to newly created console sessions after configuration.

To set current pool SSH auto mode

```
1 xe poolparam-set uuid=<pool-uuid> ssh-auto-mode=true
```

Configure automatic SSH management based on XAPI service health status. When auto-mode is enabled, SSH activation follows XAPI status: SSH is disabled when XAPI is healthy and enabled when XAPI is unhealthy.

Auto-mode enabled + `ssh_enabled_timeout = 0`: Manual SSH activation via `pool.enable_ssh` permanently disables auto-mode.

Auto-mode enabled + `ssh_enabled_timeout > 0`: Manual SSH activation temporarily disables auto-mode; auto-mode is restored to the original setting when timeout expires. (When system restart occurs during an active SSH timeout period, the original auto-mode setting is lost and is reset to enabled upon timeout/recovery.)

For the value of the pool parameter in either of these commands, you can use a pool selector. For more information, see [Pool commands](#).

You can manage the SSH access for all pools in a XenServer pool at the same time. For more information, see [Disable SSH access for a pool](#).

Pool Join / Eject Behavior

- Join –When a server joins a pool, the SSH settings (`ssh enabled timeout`, `idle console timeout`, `auto mode`) are inherited.
- Eject –When a server is ejected from a pool, the default settings are restored.
 - XenServer 8.4: SSH enabled, auto mode off

Definition of new SSH management control

Set the new XAPI API `pool.set_ssh_auto_mode` at pool level to configure SSH

`set_ssh_auto_mode`: Configure to enable or disable SSH auto-mode. When auto-mode is enabled, SSH status depends on XAPI status: SSH is disabled when XAPI is healthy, and SSH is enabled when XAPI is unhealthy.

When a user sets auto-mode to enabled and the current `ssh_enabled_timeout` is 0, enabling SSH with `pool.enable_ssh` will permanently disable auto-mode.

When a user sets auto-mode with `pool.set_auto_mode` as enabled or disabled, this is called the “original auto-mode setting.” If the current `ssh_enabled_timeout` is set to a positive number (for example, 1800 seconds), enabling SSH with `pool.enable_ssh` will temporarily disable auto-mode. When the timeout period expires, SSH will be automatically disabled and auto-mode will be restored to its original setting.

When a user sets `ssh_enabled_timeout` to a positive number (for example, 3600 seconds) and reboots the host (or restarts XAPI) before the timeout period expires, once the timeout period expires, SSH will be automatically disabled and auto-mode will be restored to true as the reboot will have lost the original auto-mode setting.

When you configure `ssh_enabled_timeout` to a positive number (for example, 180 seconds) and XAPI subsequently fails, if XAPI remains in a failed state when the timeout period expires, SSH will stay enabled until XAPI restarts. Upon XAPI restart, SSH will be automatically disabled and auto-mode will be restored to true, since the system reboot causes the loss of the original auto-mode setting.

Enable IGMP snooping on your XenServer pool

XenServer sends multicast traffic to all guest VMs leading to unnecessary load on host devices by requiring them to process packets they have not solicited. Enabling IGMP snooping prevents hosts on a local network from receiving traffic for a multicast group they have not explicitly joined, and improves the performance of multicast. IGMP snooping is especially useful for bandwidth-intensive IP multicast applications such as IPTV.

Notes:

- When enabling this feature on a pool, it may also be necessary to enable IGMP querier on one of the physical switches. Or else, multicast in the sub network falls back to broadcast and may decrease XenServer performance.
- When enabling this feature on a pool running IGMP v3, VM migration or network bond failover results in IGMP version switching to v2.
- To enable this feature with a GRE network, users must set up an IGMP Querier in the GRE network. Alternatively, you can forward the IGMP query message from the physical network into the GRE network. Or else, multicast traffic in the GRE network can be blocked.

You can enable IGMP snooping on a pool by using XenCenter or the xe CLI.

XenCenter

1. Navigate to **Pool Properties**.
2. Select **Network Options**. Here you can enable or disable IGMP snooping.

xe CLI

1. Get the pool UUID:

```
xe pool-list
```

2. Enable/disable IGMP snooping for the pool:

```
xe pool-param-set [uuid=pool-uuid] [igmp-snooping-enabled=true | false]
```

After enabling IGMP snooping, you can view the IGMP snooping table using the xe CLI.

View the IGMP snooping table

Use the following command to view the IGMP snooping table:

```
ovs-appctl mdb/show [bridge name]
```

Note:

You can get the bridge name by using `xe network-list`. These bridge names can be `xenbr0`, `xenbr1`, `xenapi`, or `xapi0`.

This outputs a table with four columns:

- port: The port of the switch (OVS).
- VLAN: The VLAN ID of the traffic.
- GROUP: The multicast group that the port solicited.
- Age: The age of this record in seconds.

If the **GROUP** is a multicast group address, this means an IGMP Report message is received on the associated switch port. This means that a receiver (member) of the multicast group is listening on this port.

Take the following example which contains two records:

port	VLAN	GROUP	Age
14	0	227.0.0.1	15
1	0	querier	24

The first record shows that there is a receiver listening on port 14 for the multicast group 227.0.0.1. The Open vSwitch forwards traffic destined for the 227.0.0.1 multicast group to listening ports for this group only (in this example, port 14), rather than broadcasting to all ports. The record linking port 14 and group 227.0.0.1 was created 15 seconds ago. By default, the timeout interval is 300 seconds. This means that if the switch does not receive any further IGMP Report messages on port 14 for 300 seconds after adding the record, the record expires and is removed from the table.

In the second record, the **GROUP** is **querier**, meaning that IGMP Query messages have been received on the associated port. A querier periodically sends IGMP Query messages, which are broadcasted to all switch ports, to determine which network nodes are listening on a multicast group. Upon receiving an IGMP Query message, the receiver responds with an IGMP Report message, which causes the receiver's multicast record to refresh and avoid expiration.

The **VLAN** column indicates to the VLAN that a receiver/querier lives. '0' means native VLAN. If you want to run multicast on some tagged VLAN, ensure that there are records on the VLAN.

Note:

For the VLAN scenario, you must have a querier record with a VLAN column value equal to the VLAN ID of the network, otherwise multicast won't work in the VLAN network.

Enable migration stream compression on your XenServer pool

During the live migration of a VM, its memory is transferred as a data stream between two hosts using the network. The migration stream compression feature compresses this data stream, speeding up the memory transfer on slow networks. This feature is disabled by default, but this can be changed by using XenCenter or the xe CLI. For more information, see [Pool Properties - Advanced](#) and [Pool parameters](#). Alternatively, you can enable compression when migrating a VM by using the command line. For more information, see the `vm-migrate` command in [VM Commands](#).

Export resource pool data

The Export Resource Data option allows you to generate a resource data report for your pool and export the report into an .xls or .csv file. This report provides detailed information about various resources in the pool such as hosts, networks, storage, virtual machines, VDIs, and GPUs. This feature enables administrators to track, plan, and assign resources based on various workloads such as CPU, storage, and network.

Note:

Export Resource Pool Data is available for XenServer Premium Edition customers.

The list of resources and various types of resource data that are included in the report:

Server:

- Name
- Pool Coordinator
- UUID
- Address
- CPU Usage
- Network (avg/max KBs)
- Used Memory
- Storage
- Uptime
- Description

Networks:

- Name
- Link Status
- MAC
- MTU
- VLAN
- Type
- Location

VDI:

- Name
- Type
- UUID
- Size
- Storage
- Description

Storage:

- Name
- Type
- UUID
- Size
- Location
- Description

VMs:

- Name
- Power State
- Running on
- Address
- MAC
- NIC
- Operating System
- Storage
- Used Memory
- CPU Usage
- UUID
- Uptime
- Template
- Description

GPU:

- Name
- Servers
- PCI Bus Path
- UUID
- Power Usage
- Temperature
- Used Memory
- Computer Utilization

Note:

Information about GPUs is available only if there are GPUs attached to your XenServer host.

To export resource data

1. In the XenCenter Navigation pane, select **Infrastructure** and then select the pool.
2. Select the **Pool** menu and then **Export Resource Data**.
3. Browse to a location where you would like to save the report and then click **Save**.

Manage your hosts

May 19, 2026

This article describes some of the actions that you can take to manage your individual XenServer® hosts.

To manage actions related to resource pools, see [Manage your pools](#).

Configure SSH access

SSH access to your XenServer hosts is enabled by default. If you want to disable SSH access to your XenServer host, run the following command:

```
1 xe host-disable-ssh host=<host_uuid_or_name_label>
```

This command disables SSH. The host refuses new SSH connections but does not disconnect any existing sessions.

To enable SSH access to your XenServer host, run the following command:

```
1 xe host-enable-ssh host=<host_uuid_or_name_label>
```

To set current host SSH temporary enabled timeout (seconds):

```
1 xe host-param-set uuid=<host-uuid> ssh-enabled-timeout=<seconds>
```

The timeout parameter determines SSH service duration. When set to 0, SSH remains active indefinitely. When set to a positive integer (seconds), SSH is disabled automatically after the specified period. Configuration changes take immediate effect on following active SSH sessions and take effect every time the user enables SSH from CLI/API.

To set current host SSH console idle timeout (seconds)

```
1 xe host-param-set uuid=<host-uuid> console-idle-timeout=<seconds>
```

Configure the idle session timeout for VNC and SSH console connections using a non-negative integer value in seconds, where 0 disables timeout (sessions never expire) and positive values (for example, 3600) automatically terminate idle sessions after the specified duration, with settings applying to newly created console sessions after configuration.

To set current host SSH auto mode

```
1 xe host-param-set uuid=<host-uuid> ssh-auto-mode=true
```

Note:

By default, the SSH service is disabled for new installations to reduce attack surface.

Configure automatic SSH management based on XAPI service health status. When auto-mode is enabled, SSH activation follows XAPI status: SSH is disabled when XAPI is healthy and enabled when XAPI is unhealthy.

Auto-mode enabled + `ssh_enabled_timeout = 0`: Manual SSH activation via `host.enable_ssh` permanently disables auto-mode.

Auto-mode enabled + `ssh_enabled_timeout > 0`: Manual SSH activation temporarily disables auto-mode; auto-mode is restored to the original setting when timeout expires. (When system restart occurs during an active SSH timeout period, the original auto-mode setting is lost and is reset to enabled upon timeout/recovery.)

For the value of the host parameter in either of these commands, you can use a host selector. For more information, see [Host selectors](#).

You can manage the SSH access for all hosts in a XenServer pool at the same time. For more information, see [Disable SSH access for a pool](#).

SSH auto-mode with answerfile

Add a new answerfile configuration `<ssh-mode>on|off|auto</ssh-mode>` to control default SSH and auto-mode settings.

- When set to 'auto', auto-mode is enabled and will automatically stop the SSH service once XAPI starts.
- When set to 'on' or 'off', auto-mode is disabled by default, and the SSH service is enabled or disabled accordingly.

Definition of new SSH management control

Set the new XAPI API `host.set_ssh_auto_mode` at host level to configure SSH

`set_ssh_auto_mode`: Configure to enable or disable SSH auto-mode. When auto-mode is enabled, SSH status depends on XAPI status: SSH is disabled when XAPI is healthy, and SSH is enabled when XAPI is unhealthy.

When a user configures auto-mode as enabled and the current `ssh_enabled_timeout` is 0, enabling SSH with `host.enable_ssh` will permanently disable auto-mode.

When a user configures auto-mode with `host.set_auto_mode` as enabled or disabled, this is called the "original auto-mode setting." If the current `ssh_enabled_timeout` is set to a positive number (for example, 1800 seconds), enabling SSH with `host.enable_ssh` will temporarily disable auto-mode. When the timeout period expires, SSH will be automatically disabled and auto-mode will be restored to its original setting.

When a user sets `ssh_enabled_timeout` to a positive number (for example, 3600 seconds) and reboots the host (or restarts XAPI) before the timeout period expires, once the timeout period expires, SSH will be automatically disabled and auto-mode will be restored to true as the reboot will have lost the original auto-mode setting.

When you configure `ssh_enabled_timeout` to a positive number (for example, 180 seconds) and XAPI subsequently fails, if XAPI remains in a failed state when the timeout period expires, SSH will stay enabled until XAPI restarts. Upon XAPI restart, SSH will be automatically disabled and auto-mode will be restored to true, since the system reboot causes the loss of the original auto-mode setting.

Install a TLS certificate on your host

The XenServer host comes installed with a default TLS certificate. However, to use HTTPS to secure communication between XenServer and Citrix Virtual Apps and Desktops™, install a certificate provided by a trusted certificate authority.

This section describes how to install certificates by using the xe CLI. For information about working with certificates by using XenCenter, see [the XenCenter documentation](#).

Ensure that your TLS certificate and its key meet the following requirements:

- The certificate and key pair are an RSA key.
- The key matches the certificate.
- The key is provided in a separate file to the certificate.
- The certificate is provided in a separate file to any intermediate certificates.
- The key file must be one of the following types: `.pem` or `.key`.
- Any certificate files must be one of the following types: `.pem`, `.cer`, or `.crt`.
- The key is greater than or equal to 2048 bits and less than or equal to 4096 bits in length.
- The key is an unencrypted PKCS #8 key and does not have a passphrase.
- The key and certificate are in base-64 encoded 'PEM' format.
- The certificate is valid and has not expired.
- The signature algorithm is SHA-256 or SHA-512.

The xe CLI warns you when the certificate and key you choose do not meet these requirements.

Where do I get a TLS certificate?

- You might already have a trusted certificate that you want to install on your XenServer host.
- Alternatively, you can create a certificate on your server and send it to your preferred certificate authority to be signed. This method is more secure as the private key can remain on the XenServer host and not be copied between systems.

Creating a TLS certificate has the following steps:

1. Generate a certificate signing request
2. Send the certificate signing request to a certificate authority
3. Install the signed certificate on your XenServer host

1. Generate a certificate signing request First, generate a private key and certificate signing request. On the XenServer host, complete the following steps:

1. To create a private key file, run the following command:

```
1 openssl genrsa -des3 -out privatekey.pem 2048
```

You are prompted for a pass phrase. This pass phrase is removed in a following step.

2. Remove the pass phrase from the key:

```
1 openssl rsa -in privatekey.pem -out privatekey.nop.pem
```


3. Create the certificate signing request by using the private key:

```
1 openssl req -new -key privatekey.nop.pem -out csr
```

4. Follow the prompts to provide the information necessary to generate the certificate signing request.

- **Country Name.** Enter the TLS Certificate country codes for your country. For example, CA for Canada or JM for Jamaica. You can find a list of TLS Certificate country codes on the web.
- **State or Province Name (full name).** Enter the state or province where the pool is located. For example, Massachusetts or Alberta.
- **Locality Name.** The name of the city where the pool is located.
- **Organization Name.** The name of your company or organization.
- **Organizational Unit Name.** Enter the department name. This field is optional.
- **Common Name.** Enter the FQDN of your XenServer host. We recommend specifying either an FQDN or an IP address that does not expire.
- **Email Address.** This email address is included in the certificate when you generate it.

The certificate signing request is saved in the current directory and is named `csr`.

5. Display the certificate signing request in the console window by running the following command:

```
1 cat csr
```

6. Copy the entire certificate signing request and use this information to request the certificate from the certificate authority.

Example certificate signing request:

```
1 -----BEGIN CERTIFICATE REQUEST-----
2 MIIDBDCCAewCAQAwYsxCzAJBgNVBAYTA1VLMRcwFQYDVQQIDA5DYW1icmlkZ2Vz
3 aGlyZTESMBAGA1UEBwwJQ2FtYnJpZGdlMRIwEAYDVQQKDA1YZW5TZXJ2ZXIxFTAT
4 ...
5 SdYCKFdo+85z8hBULFzSH6jgSP0UGQU0PcfIy7KPKyI4jnFQqeCDvLdWyhtAx9gq
6 Fu40qMSmIdNCFfnACRwYQkQgqCt/RHeUtl8srxyZC+odbunnV+ZyQdmLwLuQySUK
7 ZL8naumG3yU=
8 -----END CERTIFICATE REQUEST-----
```

2. Send the certificate signing request to a certificate authority Now that you have generated the certificate signing request, you can submit the request to your organization's preferred certificate authority.

A certificate authority (CA) is a service that provides digital certificates, you might have a CA available within your organization or, alternatively, you can use a trusted third party. Some certificate

authorities require the certificates to be hosted on a system that is accessible from the internet. We recommend not using a certificate authority with this requirement.

The certificate authority responds to your signing request and provides the following files:

- the signed certificate
- if applicable, an intermediate certificate

You can now install all these files on your XenServer host.

3. Install the signed certificate on your XenServer host After the certificate authority responds to the certificate signing request, complete the following steps to install the certificate on your XenServer host:

1. Get the signed certificate and, if the certificate authority has one, the intermediate certificate from the certificate authority.
2. Copy the key and certificates to the XenServer host.
3. Run the following command on the host:

```
1 xe host-server-certificate-install certificate=<
  path_to_certificate_file> private-key=<path_to_private_key>
  certificate-chain=<path_to_chain_file>
```

The `certificate-chain` parameter is optional.

For extra security, you can delete the private key file after the certificate is installed.

Manage the administrator password

When you first install a XenServer host, you set an administrator or *root* password. You use this password to connect XenCenter to your host or (with user name `root`) to log into **xsconsole**, the system configuration console.

If you join a host to a pool, the administrator password for the host is automatically changed to match the administrator password of the pool coordinator.

Note:

XenServer administrator passwords must contain only printable ASCII characters.

Change the password

You can use XenCenter®, the xe CLI, or **xsconsole** to change the administrator password.

XenCenter To change the administrator password for a pool or standalone host by using XenCenter, complete the following steps:

1. In the **Resources** pane, select the pool or any host in the pool.
2. On the **Pool** menu or on the **Server** menu, select **Change Server Password**.

To change the root password of a standalone host, select the host in the **Resources** pane, and click **Password** and then **Change** from the **Server** menu.

If XenCenter is configured to save your host login credentials between sessions, the new password is remembered. For more information, see [Store your host connection state](#).

After changing the administrator password, rotate the pool secret. For more information, see [Rotate the pool secret](#).

xe CLI To change the administrator password by using the xe CLI, run the following command on a host in the pool:

```
1 xe user-password-change new=<new_password>
```

Note:

Ensure that you prefix the command with a space to avoid storing the plaintext password in the command history.

After changing the administrator password, rotate the pool secret. For more information, see [Rotate the pool secret](#).

xsconsole To change the administrator password for a pool or a standalone host by using **xsconsole**, complete the following steps:

1. On the pool coordinator, go to the console.
2. Log in as `root`.
3. Type `xsconsole`. Press **Enter**. The **xsconsole** is displayed.
4. In **xsconsole**, use the arrow keys to navigate to the **Authentication** option. Press **Enter**.
5. Navigate to **Change Password**. Press **Enter**.
6. Authenticate with the administrator password.
7. In the **Change Password** dialog:
 - a) Enter your current password.
 - b) Enter a new password.

- c) Enter the new password again to confirm it.

The **Password Change Successful** screen is displayed. Press **Enter** to dismiss.

If the host is pool coordinator, this updated password is now propagated to the other hosts in the pool.

After changing the administrator password, rotate the pool secret. For more information, see [Rotate the pool secret](#).

Reset a lost root password

If you lose the administrator (root) password for your XenServer host, you can reset the password by accessing the host directly.

1. Reboot the XenServer host.
2. When the GRUB menu shows, press **e** to edit the boot menu entry.
3. Add `init=/sysroot/bin/sh` to the line that starts with `module2`.
4. Press **Ctrl-X** to boot into a root shell.
5. At the command shell, run the following commands:

```
1 chroot /sysroot
2 passwd
3
4 (type the new password twice)
5
6 sync
7 /sbin/reboot -f
```

If the host is pool coordinator, this updated password is now propagated to the other hosts in the pool.

After changing the administrator password, rotate the pool secret. For more information, see [Rotate the pool secret](#)

Change the NTP configuration on a server

You can update the NTP configuration for your server from xsconsole.

1. In the host console, type `xsconsole`.
2. In xsconsole, go to **Network and Management Interfaces > Network Time (NTP)**.
3. Enter your password to proceed.
4. From the **Configure Network Time** menu, choose the option you want to configure.

You can also configure NTP using the **XenAPI**.

This allows you to automate NTP configuration without directly accessing dom0 through SSH.

Example using Python SDK:

```
session.xenapi.host.set_ntp_mode(host_ref, 'Factory')
```

NTP modes:

- Factory - uses built-in XenServer NTP pool servers
- DHCP - uses DHCP-assigned NTP servers
- Custom - uses NTP servers configured by the user
- Disabled - disables NTP (not recommended)

To check NTP sync status:

- `session.xenapi.host.get_ntp_synchronized(host_ref)`
- `session.xenapi.host.get_ntp_servers_status(host_ref)`

Note:

If NTP is configured directly on dom0 instead of through the XenAPI, values may become inconsistent with XenAPI fields.

To resynchronize the values with XAPI, restart the toolstack:

```
xe-toolstack-restart
```

Change the timezone on a server

You can configure the timezone for the host using the **XenAPI**.

Example using Python SDK:

```
session.xenapi.host.set_timezone(host_ref, 'UTC')
```

Timezone values must use IANA timezone names, for example:

- UTC
- Europe/Paris
- Asia/Kolkata

To read the current timezone:

```
session.xenapi.host.get_timezone(host_ref)
```

To list supported timezones:

```
session.xenapi.host.list_timezones(host_ref)
```

Note:

If Timezone is configured directly on dom0 instead of through the XenAPI, values may become inconsistent with XenAPI fields.

To resynchronize the values with XAPI, restart the toolstack:

```
xe-toolstack-restart
```

Change max-cstate on a server

The max-cstate setting controls the deepest CPU idle power state allowed.

Supported values:

- "" - unlimited (default)
- 0 - limit to C0
- 1 - limit to C1

Example using Python SDK:

```
session.xenapi.host.set_max_cstate(host_ref, '1')
```

To read the current value:

```
session.xenapi.host.get_max_cstate(host_ref)
```

Note:

If max-cstate is configured directly on dom0 instead of through the XenAPI, values may become inconsistent with XenAPI fields.

To resynchronize the values with XAPI, restart the toolstack:

```
xe-toolstack-restart
```

Prepare a pool of XenServer hosts for maintenance

Before performing maintenance operations on a host that is part of a resource pool, you must disable it. Disabling the host prevents any VMs from being started on it. You must then migrate its VMs to another XenServer host in the pool. You can do this by placing the XenServer host in to Maintenance mode using XenCenter. For more information, see [Run in maintenance mode](#) in the XenCenter documentation.

Backup synchronization occurs every 24 hrs. Placing the pool coordinator in maintenance mode results in the loss of the last 24 hrs of RRD updates for offline VMs.

Warning:

We highly recommend rebooting all XenServer hosts before installing an update and then verifying their configuration. Some configuration changes only take effect when the XenServer host is rebooted, so the reboot might uncover configuration problems that can cause the update to fail.

To prepare a host in a pool for maintenance operations by using the CLI

1. Run the following command:

```
1 xe host-disable uuid=XenServer_host_uuid
2 xe host-evacuate uuid=XenServer_host_uuid
```

This command disables the XenServer host and then migrates any running VMs to other XenServer hosts in the pool.

2. Perform the desired maintenance operation.
3. Enable the XenServer host when the maintenance operation is complete:

```
1 xe host-enable
```

4. Restart any halted VMs and resume any suspended VMs.

Shut down or restart a XenServer host

We recommend always using the xe CLI or XenCenter to shut down or restart your XenServer hosts. Do not use other methods (such as hardware-based resets or operating system commands) to shut down or restart your hosts during normal operation. Doing so can cause unexpected behaviors on your host and pool.

To shut down or restart a host by using the xe CLI, use the following commands:

1. Prevent any new VMs from starting on or being migrated to the specified host:

```
1 xe host-disable host=<host name>
```

2. If required, migrate any running VMs off the host:

```
1 xe host-evacuate uuid=<host uuid>
```

3. Shut down or restart the host with one of the following xe CLI commands.

- To cleanly shut down the host:

““

xe host-shutdown host=

- To cleanly restart the host:

```
1  xe host-reboot host=<host name>
```

For how to shut down a server by using XenCenter, see the [XenCenter documentation](#).

Configure host power-on

Powering on hosts remotely

You can use the XenServer host Power On feature to turn a host on and off remotely, either from XenCenter or by using the CLI.

To enable host power, the host must have one of the following power-control solutions:

- **Wake on LAN enabled network card.**
- **Intelligent Platform Management Interface (IPMI).**
- A custom script based on the management API that enables you to turn the power on and off through XenServer. For more information, see *Configuring a custom script for the Host Power On feature* in the following section.

Using the Host Power On feature requires two tasks:

1. Ensure the hosts in the pool support controlling the power remotely. For example, they have Wake on LAN functionality or support IPMI, or you have created a custom script.
2. Enable the Host Power On functionality using the CLI or XenCenter.

Use the CLI to manage host power-on

You can manage the Host Power On feature using either the CLI or XenCenter. This section provides information about managing it with the CLI.

Host Power On is enabled at the host level (that is, on each XenServer).

After you enable Host Power On, you can turn on hosts using either the CLI or XenCenter.

To enable host power-on by using the CLI Run the command:

```
1  xe host-set-power-on-mode host=<host uuid> \  
2    power-on-mode=("wake-on-lan", "IPMI", "custom") \  
3    power-on-config=key:value
```


To turn on hosts remotely by using the CLI Run the command:

```
1 xe host-power-on host=<host uuid>
```

Certificate verification

March 3, 2026

When certificate verification is enabled for a pool, all TLS communication endpoints on its management network use certificates to validate the identity of their peers before transmitting confidential information.

Behavior

Connections initiated by a XenServer® host on the management network require the destination endpoint to provide a TLS certificate to verify its identity. This requirement affects the following items that are part of the pool or interact with the pool:

- Hosts in the pool
- XenCenter®
- Third-party clients that use the API

Certificate verification is compatible with both the self-signed certificates provided by XenServer and user-installed certificates signed by a trusted authority. For more information, see [Install a TLS certificate on your host](#).

Each XenServer host in a pool has two certificates that identify it:

- *Pool-internal identity certificates* are used to secure communications between hosts within the pool. For communication within the pool, XenServer always uses self-signed certificates.
- *Server identity certificates* are used to verify the identity of a XenServer host to any client applications that communicate with the pool on the management network. For communication between the host and a client application, you can use self-signed certificates or you can install your own TLS certificates on your hosts.

When a host first joins the pool or a client first makes a connection to the pool, the pool trusts the connection. During this first connection, certificates are exchanged between the pool and the joining host or the connecting client. For all subsequent communications by this host or client on the management network, the certificates are used to verify the identity of the parties involved in the communication.

We recommend that you enable certificate verification on all your hosts and pools. For a XenServer host to successfully join a pool, both the host and the pool must have certificate verification either enabled or disabled. If certificate verification is enabled on one and not the other, the join operation is not successful. XenCenter provides a warning message that advises you to enable certificate verification on the pool or on the joining host.

When a host leaves a pool with certificate verification enabled, both the host and the pool delete the certificates that relate to the other.

The Workload Balancing virtual appliance can be used with certificate verification. You must ensure that the Workload Balancing self-signed certificates are installed into your XenServer host.

The Conversion Manager virtual appliance does not connect to XenServer hosts and so is exempt from the certification checking requirement when it acts as a TLS client end point.

Enabling certificate verification for your pool

Certificate verification is enabled by default on fresh installations of XenServer 8.4 and later. If you upgrade from an earlier version of XenServer, certificate verification is not enabled automatically and you must enable it. XenCenter prompts you to enable certificate verification the next time you connect to the upgraded pool.

Before enabling certificate verification on a pool, ensure that no operations are running in the pool.

Enable by using XenCenter

XenCenter provides several ways to enable certificate verification.

- When first connecting XenCenter to a pool without certificate verification enabled, you are prompted to enable it. Click **Yes, Enable certificate verification**.
- In the **Pool** menu, select **Enable Certificate Verification**.
- On the **General** tab of the pool, right-click the entry **Certificate Verification** and choose **Enable Certificate Verification** from the menu.

Enable by using the xe CLI

To enable certificate verification for a pool, run the following command in the console of a host in the pool:

```
1 xe pool-enable-tls-verification
```

Managing certificates

You can install, view information about, and reset the certificates that are used to verify the identity of a host.

Installing certificates

You can install your own TLS certificate for the host to present as its identity certificate when receiving connections from client applications on the management network.

For more information, see [Install a TLS certificate on your host](#).

Viewing certificate information

To find out whether a pool has certificate verification enabled:

- In XenCenter, look in the **General** tab for the pool. The **General** section has an entry for **Certificate Verification** which shows whether certificate verification is enabled or disabled. This tab also contains a **Certificates** section that lists the name, validity, and thumbprint for the CA certificates.
- With the xe CLI, you can run the following command:

```
1 xe pool-param-get uuid=<pool_uuid> param-name=tls-verification-enabled
```

If certificate verification is enabled, the line `tls-verification-enabled ( R0 ) : true` appears in the command output.

To view information about the certificates on a XenServer host:

- In XenCenter, go to the **General** tab for that host. The **Certificates** section shows the thumbprint and the validity dates for the server identity certificate and the pool-internal identity certificate.
- With the xe CLI, you can run the following command:

```
1 xe certificate-list
```

Refreshing pool-internal identity certificates

You can refresh the pool-internal identity certificate by using the xe CLI:

1. Find the UUID of the host whose certificate you want to reset by running the following command:

```
1 xe host-list
```

2. To reset the certificate, run the following command:

```
1 xe host-refresh-server-certificate host=<host_uuid>
```

Note:

Any host selector parameter can be used with this command to indicate the host to reset the certificate on.

Resetting server identity certificates

You can reset the server identity certificate from XenCenter or the xe CLI. Resetting a certificate deletes the certificate from the host and installs a new self-signed certificate in its place.

To reset a certificate in XenCenter:

1. Go to the **General** tab for the host.
2. In the **Certificates** section, right-click on the certificate you want to reset.
3. From the menu, select **Reset Certificate**.
4. In the dialog that appears, click **Yes** to confirm the certificate reset.

Alternatively, in the **Server** menu, you can go to **Certificates > Reset Certificate**.

When you reset a certificate, any existing connections to the XenServer host are disconnected—including the connection between XenCenter and the host. XenCenter reconnects automatically to the host after a certificate reset.

To reset a certificate by using the xe CLI:

1. Find the UUID of the host whose certificate you want to reset by running the following command:

```
1 xe host-list
```

2. To reset the certificate, run the following command:

```
1 xe host-reset-server-certificate host=<host_uuid>
```

Note:

Any host selector parameter can be used with this command to indicate the XenServer host to reset the certificate on.

When you reset a certificate, any existing connections to the XenServer host are disconnected—including the connection between XenCenter and the host. XenCenter reconnects automatically to the host after a certificate reset.

Expiry alerts

XenCenter shows alerts in the **Notifications** view when your server identity certificates, pool-internal identity certificates, or pool CA certificates are close to their expiry date.

Temporarily disabling certificate verification

We do not recommend that you disable certificate verification after it has been enabled on a host or pool. However, XenServer provides commands that can be used to disable certificate verification on a per host basis when troubleshooting problems with certificates.

To temporarily disable certificate verification, run the following command on the host console:

```
1 xe host-emergency-disable-tls-verification
```

XenCenter shows an alert in the **Notifications** view when certificate verification is disabled on a host in a pool where the feature is enabled.

After you have resolved any issues with certificates on the host, ensure that you enable certificate verification on it again. To enable certificate verification again, run the following command on the host console:

```
1 xe host-emergency-reenable-tls-verification
```

GFS2 clustered pools

April 27, 2026

GFS2 clustering provides extra features that are required for resource pools that use GFS2 SRs. For more information about GFS2, see [Configure storage](#).

A GFS2 cluster is a pool of up to 16 XenServer® hosts that are more closely connected and coordinated than hosts in non-clustered pools. The hosts in the GFS2 cluster maintain constant communication with each other on a selected network. All hosts in the GFS2 cluster are aware of the state of every host in the GFS2 cluster. This host coordination enables the GFS2 cluster to control access to the contents of the GFS2 SR.

Note:

The GFS2 clustering feature only benefits pools that contain a GFS2 SR. If your pool does not contain a GFS2 SR, do not enable GFS2 clustering in your pool.

Quorum

Each host in a GFS2 cluster must always be in communication with the majority of hosts in the GFS2 cluster (including itself). This state is known as a host having quorum. If a host does not have quorum, that host self-fences.

The number of hosts that must be in communication to initially achieve quorum can be different to the number of hosts a GFS2 cluster requires to keep quorum.

The following table summarizes this behavior. The value of *n* is the total number of hosts in the GFS2 clustered pool.

	Number of hosts required to achieve quorum	Number of hosts required to remain quorate
Odd number of hosts in the pool	$(n+1)/2$	$(n+1)/2$
Even number of hosts in the pool	$(n/2)+1$	$n/2$

For a GFS2 clustered pool, you can check if the pool has quorum by querying the `is-quorate` parameter of the GFS2 cluster:

```
1 xe cluster-list params=is-quorate uuid=<cluster_id>
```

To see how many hosts in the GFS2 cluster are live, run the following command:

```
1 xe cluster-list params=live-hosts uuid=<cluster_id>
```

To see how many live hosts are required for the GFS2 cluster to achieve quorum, run the following command:

```
1 xe cluster-list params=quorum uuid=<cluster_id>
```

When the GFS2 cluster is created, the number of live hosts must be greater than or equal to this value. To retain quorum, the required number of hosts can be different to the value returned by this command depending on whether the GFS2 cluster contains an odd or even number of hosts.

Odd-numbered pools

To achieve the quorum value for an odd-numbered pool you require half of one more than the total number of hosts in the GFS2 cluster: $(n+1)/2$. This is also the minimum number of hosts that must remain contactable for the pool to remain quorate.

For example, in a 5-host GFS2 clustered pool, 3 hosts must be contactable for the GFS2 cluster to both become active and remain quorate $[(5+1)/2 = 3]$.

Where possible it is recommended to use an odd number of hosts in a GFS2 clustered pool as this ensures that hosts are always able to determine if they have a quorate set.

Even-numbered pools

When an even-numbered GFS2 clustered pool powers up from a cold start, $(n/2)+1$ hosts must be available before the hosts have quorum. After the hosts have quorum, the GFS2 cluster becomes active.

However, an active even-numbered pool can remain quorate if the number of contactable hosts is at least $n/2$. As a result, it is possible for a running GFS2 cluster with an even number of hosts to split exactly in half. The running GFS2 cluster decides which half of the GFS2 cluster self-fences and which half of the GFS2 cluster has quorum. The half of the GFS2 cluster that contains the node with the lowest ID that was seen as active before the GFS2 cluster split remains active and the other half of the GFS2 cluster self-fences.

For example, in a 4-host GFS2 clustered pool, 3 hosts must be contactable for the GFS2 cluster to become active $[4/2 + 1 = 3]$. After the GFS2 cluster is active, to remain quorate, only 2 hosts must be contactable $[4/2 = 2]$ and that set of hosts must include the host with the lowest node ID known to be active.

Self-fencing

If a host detects that it does not have quorum, it self-fences within a few seconds. When a host self-fences, it restarts immediately. All VMs running on the host are immediately stopped because the host does a hard shutdown. In a GFS2 clustered pool that uses high availability, XenServer restarts the VMs according to their restart configuration on other pool members. The host that self-fenced restarts and attempts to rejoin the GFS2 cluster.

If the number of live hosts in the GFS2 cluster becomes less than the quorum value, all the remaining hosts lose quorum.

In an ideal scenario, your GFS2 clustered pool always has more live hosts than are required for quorum and XenServer never fences. To make this scenario more likely, consider the following recommendations when setting up your GFS2 clustered pool:

- Ensure that you have good hardware redundancy.
- Use a dedicated bonded network for the GFS2 cluster network. Ensure that the bonded NICs are on the same L2 segment. For more information, see [Networking](#).
- Configure storage multipathing between the pool and the GFS2 SR. For more information, see [Storage multipathing](#).

Create a GFS2 clustered pool

Before you begin, ensure the following prerequisites are met:

- All XenServer hosts in the GFS2 clustered pool must have at least 2 GiB of control domain memory.

Depending on your environment, your hosts might require more control domain memory than this. If you have insufficient control domain memory on your hosts, your pool can experience network instability. Network instability can cause problems for a GFS2 clustered pool with GFS2 SRs. For information about changing the amount of control domain memory and monitoring the memory behavior, see [Memory usage](#).

- All hosts in the GFS2 cluster must use static IP addresses for the GFS2 cluster network.
- We recommend that you use GFS2 clustering only in pools containing at least three hosts, as pools of two hosts are sensitive to self-fencing the entire pool.
- GFS2 clustered pools only support up to 16 hosts per pool.
- If you have a firewall between the hosts in your pool, ensure that hosts can communicate on the GFS2 cluster network using the following ports:
 - TCP: 8892, 8896, 21064
 - UDP: 5404, 5405

For more information, see [Communication ports used by XenServer](#).

- If you are adding GFS2 clustering to an existing pool, ensure that high availability is disabled. You can enable high availability again after GFS2 clustering is enabled.
- We strongly recommend that you use a bonded network for your GFS2 clustered pool that is not used for any other traffic.

If you prefer, you can set up GFS2 clustering on your pool by using XenCenter. For more information, see the [XenCenter product documentation](#).

To use the xe CLI to create a GFS2 clustered pool:

1. Create a bonded network to use as the GFS2 cluster network.

Note:

We strongly recommend that you use a dedicated bonded network for your GFS2 clustered pool. Do not use this network for any other traffic.

On the XenServer host that you want to be the pool coordinator, complete the following steps:

- a) Open a console on the XenServer host.

- b) Create a network for use with the bonded NIC by using the following command:

```
1 xe network-create name=label=bond0
```

The UUID of the new network is returned.

- c) Find the UUIDs of the PIFs to use in the bond by using the following command:

```
1 xe pif-list
```

- d) Create your bonded network in either active-active mode, active-passive mode, or LACP bond mode. Depending on the bond mode you want to use, complete one of the following actions:

- To configure the bond in active-active mode (default), use the `bond-create` command to create the bond. Using commas to separate the parameters, specify the newly created network UUID and the UUIDs of the PIFs to be bonded:

```
1 xe bond-create network-uuid=<network_uuid> /  
2     pif-uuids=<pif_uuid_1>,<pif_uuid_2>,<pif_uuid_3>,<  
    pif_uuid_4>
```

Type two UUIDs when you are bonding two NICs and four UUIDs when you are bonding four NICs. The UUID for the bond is returned after running the command.

- To configure the bond in active-passive or LACP bond mode, use the same syntax, add the optional `mode` parameter, and specify `lacp` or `active-backup`:

```
1 xe bond-create network-uuid=<network_uuid> pif-uuids=<  
    pif_uuid_1>, /  
2     <pif_uuid_2>,<pif_uuid_3>,<pif_uuid_4> /  
3     mode=balance-slb | active-backup | lacp
```

After you have created your bonded network on the pool coordinator, when you join other XenServer hosts to the pool, the network and bond information is automatically replicated to the joining server.

For more information, see [Networking](#).

2. Create a resource pool of at least three XenServer hosts.

Repeat the following steps on each XenServer host that is a (non-master) pool member:

- a) Open a console on the XenServer host.
- b) Join the XenServer host to the pool on the pool coordinator by using the following command:

```
1 xe pool-join master-address=master_address master-username=  
    administrators_username master-password=password
```

The value of the `master-address` parameter must be set to the fully qualified domain name of the XenServer host that is the pool coordinator. The `password` must be the administrator password set when the pool coordinator was installed.

For more information, see [Hosts and resource pools](#).

3. For every PIF that belongs to this network, set `disallow-unplug=true`.

- a) Find the UUIDs of the PIFs that belong to the network by using the following command:

```
1 xe pif-list
```

- b) Run the following command on a XenServer host in your resource pool:

```
1 xe pif-param-set disallow-unplug=true uuid=<pif_uuid>
```

4. Enable GFS2 clustering on your pool. Run the following command on a XenServer host in your resource pool:

```
1 xe cluster-pool-create network-uuid=<network_uuid>
```

Provide the UUID of the bonded network that you created in an earlier step.

Disable GFS2 clustering

You can disable GFS2 clustering. After you disable GFS2 clustering, the pool continues to exist, but is no longer GFS2 clustered and can no longer use GFS2 SRs.

To disable GFS2 clustering, run the following command:

```
1 xe cluster-pool-destroy cluster-uuid=<uuid>
```

Manage your GFS2 clustered pool

When managing your GFS2 clustered pool, the following practices can decrease the risk of the pool losing quorum.

Add or remove a host on a GFS2 clustered pool

When adding or removing a host on a GFS2 clustered pool, ensure that all the hosts in the GFS2 cluster are online.

You can add or remove a host on a GFS2 clustered pool by using XenCenter. For more information, see [Add a Server to a Pool](#) and [Remove a Server From a Pool](#).

You can also add or remove a host on a GFS2 clustered pool by using the `xe` CLI. For more information, see [Add a host to a pool by using the `xe` CLI](#) and [Remove XenServer hosts from a resource pool](#).

Ensure that hosts are shut down cleanly

When a host is cleanly shut down, it is temporarily removed from the GFS2 cluster until it is started again. While the host is shut down, it does not count toward the quorum of the GFS2 cluster. The host's absence does not cause other hosts to lose quorum. For more information, see [Shut down a XenServer host](#).

However, if a host is forcibly or unexpectedly shut down, it is not removed from the GFS2 cluster before it goes offline. This host does count toward the quorum value of the GFS2 cluster. Its shutdown can cause other hosts to lose quorum.

If it is necessary to shut down a host forcibly, first check how many live hosts are in the GFS2 cluster. You can do this with the command `corosync-quorumtool`. In the command output, the number of live hosts is the value of `Total votes:` and the number of live hosts required to retain quorum is the value of `Quorum:`.

- If the number of live hosts is the same as the number of hosts needed to remain quorate, do not forcibly shut down the host. Doing so causes the whole GFS2 cluster to fence.

Instead, attempt to recover other hosts and increase the live hosts number before forcibly shutting down the host.

- If the number of live hosts is close to the number of hosts needed to remain quorate, you can forcibly shut down the host. However, this makes the GFS2 cluster more vulnerable to fully fencing if other hosts in the pool have issues.

Always try to restart the shut down host as soon as possible to increase the resiliency of your GFS2 cluster.

Use maintenance mode

Before doing something on a host that might cause that host to lose quorum, put the host into maintenance mode. When a host is in maintenance mode, running VMs are migrated off it to another host in the pool. Also, if that host was the pool coordinator, that role is passed to a different host in the pool. If your actions cause a host in maintenance mode to self-fence, you don't lose any VMs or lose your XenCenter® connection to the pool.

Hosts in maintenance mode still count towards the quorum value for the GFS2 cluster.

You can only change the IP address of a host that is part of a GFS2 clustered pool when that host is in maintenance mode. Changing the IP address of a host causes the host to leave the GFS2 cluster.

When the IP address has been successfully changed, the host rejoins the GFS2 cluster. After the host rejoins the GFS2 cluster, you can take it out of maintenance mode.

Recover hosts that have self-fenced or are offline

It is important to recover hosts that have self-fenced. While these GFS2 cluster members are offline, they count towards the quorum number for the GFS2 cluster and decrease the number of GFS2 cluster members that are contactable. This situation increases the risk of a subsequent host failure causing the GFS2 cluster to lose quorum and shut down completely.

Having offline hosts in your GFS2 cluster also prevents you from performing certain actions. In a GFS2 clustered pool, every member of the pool must agree to every change of pool membership before the change can be successful. If a GFS2 cluster member is not contactable, XenServer prevents operations that change GFS2 cluster membership (such as host add or host remove).

Mark hosts as unrecoverable

If one or more offline hosts cannot be recovered, you can tell the GFS2 clustered pool to forget them. These hosts are permanently removed from the pool. After hosts are removed from the GFS2 clustered pool, they no longer count towards the quorum value.

To mark a host as unrecoverable, use the following command:

```
1 xe host-forget uuid=<host_uuid>
```

Recover a forgotten host

After a GFS2 clustered pool is told to forget a host, the host cannot be added back into the pool.

To rejoin the GFS2 clustered pool, you must reinstall XenServer on the host so that it appears as a new host to the pool. You can then join the host to the GFS2 clustered pool in the usual way.

Troubleshoot your GFS2 clustered pool

If you encounter issues with your GFS2 clustered pool, see [Troubleshoot GFS2 clustered pools](#).

Constraints

- GFS2 clustered pools only support up to 16 hosts per pool.
- To enable HA on your GFS2 clustered pool, the heartbeat SR must be a GFS2 SR.

- For GFS2 cluster traffic, we strongly recommend that you use a bonded network that uses at least two different network switches. Do not use this network for any other purposes.
- Changing the IP address of the GFS2 cluster network by using XenCenter requires GFS2 clustering and all GFS2 SRs to be temporarily disabled.
- Do not change the bonding of your GFS2 cluster network while the GFS2 cluster is live and has running VMs. This action can cause hosts in the GFS2 cluster to hard restart (fence).
- If you have an IP address conflict (multiple hosts having the same IP address) on your GFS2 cluster network involving at least one host with GFS2 clustering enabled, the GFS2 cluster does not form correctly and the hosts are unable to fence when required. To fix this issue, resolve the IP address conflict.

Troubleshoot GFS2 clustered pools

March 16, 2026

XenServer® pools that use GFS2 to thin provision their shared block storage are clustered. These pools behave differently to pools that use shared file-based storage or LVM with shared block storage. As a result, there are some specific issues that might occur in XenServer clustered pools and GFS2 environments.

Use the following information to troubleshoot minor issues that might occur when using this feature.

All my hosts can ping each other, but I can't create a cluster. Why?

The clustering mechanism uses specific ports. If your hosts can't communicate on these ports (even if they can communicate on other ports), you can't enable clustering for the pool.

Ensure that the hosts in the pool can communicate on the following ports:

- TCP: 8892, 8896, 21064
- UDP: 5404, 5405 (not multicast)

If there are any firewalls or similar between the hosts in the pool, ensure that these ports are open.

If you have previously configured HA in the pool, disable the HA before enabling clustering.

Why am I getting an error when I try to join a new host to an existing clustered pool?

When clustering is enabled on a pool, every pool membership change must be agreed by every member of the cluster before it can succeed. If a cluster member isn't contactable, operations that change

cluster membership (such as host add or host remove) fail.

To add your new host to the clustered pool:

1. Ensure that all of your hosts are online and can be contacted.
2. Ensure that the hosts in the pool can communicate on the following ports:
 - TCP: 8892, 8896, 21064
 - UDP: 5404, 5405 (not multicast)
3. Ensure that the joining host has an IP address allocated on the NIC that joins the cluster network of the pool.
4. Ensure that clustering is not using a non-management VLAN network in the pool.
5. Ensure that no host in the pool is offline when a new host is trying to join the clustered pool.
6. If an offline host cannot be recovered, mark it as dead to remove it from the cluster. For more information, see [A host in my clustered pool is offline and I can't recover it. How do I remove the host from my cluster?](#)

What do I do if some members of the clustered pool aren't joining the cluster automatically?

This issue might be caused by members of the clustered pool losing synchronization.

To resync the members of the clustered pool, use the following command:

```
1 xe cluster-pool-resync cluster-uuid=<cluster_uuid>
```

If the issue persists, you can try to reattach the GFS2 SR. You can do this task by using the xe CLI or through XenCenter.

Reattach the GFS2 SR by using the xe CLI:

1. Detach the GFS2 SR from the pool. On each host, run the xe CLI command `xe pbd-unplug uuid=<uuid_of_pbd>`.
2. Disable the clustered pool by using the command `xe cluster-pool-destroy cluster-uuid=<cluster_uuid>`

If the preceding command is unsuccessful, you can forcibly disable a clustered pool by running `xe cluster-host-force-destroy uuid=<cluster_host>` on every host in the pool.

3. Enable the clustered pool again by using the command `xe cluster-pool-create network-uuid=<network_uuid> [cluster-stack=cluster_stack] [token-timeout=token_timeout] [token-timeout-coefficient=token_timeout_coefficient]`

4. Reattach the GFS2 SR by running the command `xe pbd-plug uuid=<uuid_of_pbd>` on each host.

Alternatively, to use XenCenter to reattach the GFS2 SR:

1. In the pool **Storage** tab, right-click on the GFS2 SR and select **Detach....**
2. From the toolbar, select **Pool > Properties**.
3. In the **Clustering** tab, deselect **Enable clustering**.
4. Click **OK** to apply your change.
5. From the toolbar, select **Pool > Properties**.
6. In the **Clustering** tab, select **Enable clustering** and choose the network to use for clustering.
7. Click **OK** to apply your change.
8. In the pool **Storage** tab, right-click on the GFS2 SR and select **Repair**.

How do I know if my host has self-fenced?

If your host self-fenced, it might have rejoined the cluster when it restarted. To see if a host has self-fenced and recovered, you can check the `/var/opt/xapi-clusterd/boot-times` file to see the times the host started. If there are start times in the file that you did not expect to see, the host has self-fenced.

Why is my host offline? How can I recover it?

There are many possible reasons for a host to go offline. Depending on the reason, the host can either be recovered or not.

The following reasons for a host to be offline are more common and can be addressed by recovering the host:

- Clean shutdown
- Forced shutdown
- Temporary power failure
- Reboot

The following reasons for a host to be offline are less common:

- Permanent host hardware failure
- Permanent host power supply failure
- Network partition
- Network switch failure

These issues can be addressed by replacing hardware or by marking failed hosts as dead.

A host in my clustered pool is offline and I can't recover it. How do I remove the host from my cluster?

You can tell the cluster to forget the host. This action removes the host from the cluster permanently and decreases the number of live hosts required for quorum.

To remove an unrecoverable host, use the following command:

```
1 xe host-forget uuid=<host_uuid>
```

This command removes the host from the cluster permanently and decreases the number of live hosts required for quorum.

Note:

If the host isn't offline, this command can cause data loss. You're asked to confirm that you're sure before proceeding with the command.

After a host is forgotten, it can't be added back into the cluster. To add this host back into the cluster, you must do a fresh installation of XenServer on the host.

I've repaired a host that was marked as dead. How do I add it back into my cluster?

A XenServer host that has been marked as dead can't be added back into the cluster. To add this system back into the cluster, you must do a fresh installation of XenServer. This fresh installation appears to the cluster as a new host.

What do I do if my cluster keeps losing quorum and its hosts keep fencing?

If one or more of the XenServer hosts in the cluster gets into a fence loop because of continuously losing quorum, you can boot the host with the `nocluster` dom0 command-line argument. Connect to the physical or serial console of the host and at the dom0 command line, run the following command: `/opt/xensource/libexec/xen-cmdline --set-dom0 nocluster`. Next time the host reboots, it does not attempt to join the cluster.

After diagnosing and resolving the issue with the host, you can enable clustering by removing the `nocluster` argument. To do so, run the following command: `/opt/xensource/libexec/xen-cmdline --remove-dom0 nocluster`. Reboot your host to make the change take effect.

For more information about accessing hosts through the serial console and editing the Xen command line, see [Advanced troubleshooting](#).

What happens when the pool coordinator gets restarted in a clustered pool?

In most cases, the behavior when the pool coordinator is shut down or restarted in a clustered pool is the same as that when another pool member shuts down or restarts.

How the host is shut down or restarted can affect the quorum of the clustered pool. For more information about quorum, see [Quorum](#).

The only difference in behavior depends on whether HA is enabled in your pool:

- If HA is enabled, a new coordinator is selected and general service is maintained.
- If HA is not enabled, there is no coordinator for the pool. Running VMs on the remaining hosts continue to run. Most administrative operations aren't available until the coordinator restarts.

Why has my pool vanished after a host in the clustered pool is forced to shut down?

If you shut down a host normally (not forcibly), it's temporarily removed from quorum calculations until it's turned back on. However, if you forcibly shut down a host or it loses power, that host still counts towards quorum calculations. For example, if you had a pool of 3 hosts and forcibly shut down 2 of them the remaining host fences because it no longer has quorum.

Try to always shut down hosts in a clustered pool cleanly. For more information, see [Manage your clustered pool](#).

Why did all hosts within the clustered pool restart at the same time?

All hosts in an active cluster are considered to have lost quorum when the number of contactable hosts in the pool is less than these values:

- For a pool with an even number of hosts: $n/2$
- For a pool with an odd number of hosts: $(n+1)/2$

The letter n indicates the total number of hosts in the clustered pool. For more information about quorum, see [Quorum](#).

In this situation, all hosts self-fence and you see all hosts restarting.

To diagnose why the pool lost quorum, the following information can be useful:

- In XenCenter, check the **Notifications** section for the time of the issue to see whether self-fencing occurred.
- On the cluster hosts, check `/var/opt/xapi-clusterd/boot-times` to see whether a reboot occurred at an unexpected time.
- In `Crit.log`, check whether any self-fencing messages are outputted.

- Review the `dlm_tool status` command output for fencing information.

Example `dlm_tool status` output:

```
1 dlm_tool status
2
3 cluster nodeid 1 quorate 1 ring seq 8 8
4 daemon now 4281 fence_pid 0
5 node 1 M add 3063 rem 0 fail 0 fence 0 at 0 0
6 node 2 M add 3066 rem 0 fail 0 fence 0 at 0 0
```

When collecting logs for debugging, collect diagnostic information from all hosts in the cluster. In the case where a single host has self-fenced, the other hosts in the cluster are more likely to have useful information.

Collect full server status reports for the hosts in your clustered pool. For more information, see [XenServer server status reports](#).

Why can't I recover my clustered pool when I have quorum?

If you have a clustered pool with an even number of hosts, the number of hosts required to *achieve* quorum is one more than the number of hosts required to *retain* quorum. For more information about quorum, see [Quorum](#).

If you are in an even-numbered pool and have recovered half of the hosts, you must recover one more host before you can recover the cluster.

You can check whether your cluster has quorum by running the following command:

```
1 xe cluster-list params=is-quorate uuid=<cluster_id>
```

Why do I see an Invalid token error when changing the cluster settings?

When updating the configuration of your cluster, you might receive the following error message about an invalid token (`"[["InternalError\\",\\"Invalid token\\"]]"`).

You can resolve this issue by completing the following steps:

1. (Optional) Back up the current cluster configuration by collecting a server status report that includes the xapi-clusterd and system logs.
2. Use XenCenter to detach the GFS2 SR from the clustered pool.
In the pool **Storage** tab, right-click on the GFS2 SR and select **Detach....**
3. On any host in the cluster, run this command to forcibly destroy the cluster:

```
1 xe cluster-pool-force-destroy cluster-uuid=<uuid>
```

4. Use XenCenter to reenable clustering on your pool.
 - a) From the toolbar, select **Pool > Properties**.
 - b) In the **Clustering** tab, select **Enable clustering** and choose the network to use for clustering.
 - c) Click **OK** to apply your change
5. Use XenCenter to reattach the GFS2 SR to the pool

In the pool **Storage** tab, right-click on the GFS2 SR and select **Repair**.

Manage users

June 23, 2026

Defining users, groups, roles and permissions allows you to control who has access to your XenServer® hosts and pools and what actions they can perform.

When you first install XenServer, a user account is added to XenServer automatically. This account is the local super user (LSU), or root, which XenServer authenticates locally.

The LSU, or root, is a special user account intended for system administration and has all permissions. In XenServer, the LSU is the default account at installation. XenServer authenticates the LSU account. LSU does not require any external authentication service. If an external authentication service fails, the LSU can still log in and manage the system. The LSU can always access the XenServer physical server through SSH.

You can create more users by adding the Active Directory accounts through either XenCenter's Users tab or the xe CLI. If your environment does not use Active Directory, you are limited to the LSU account.

Note:

When you create users, XenServer does not assign newly created user accounts RBAC roles automatically. Therefore, these accounts do not have any access to the XenServer pool until you assign them a role.

These permissions are granted through roles, as discussed in the *Authenticating users with Active Directory (AD)* section.

Authenticate users with Active Directory (AD)

If you want to have multiple user accounts on a host or a pool, you must use Active Directory user accounts for authentication. AD accounts let XenServer users log on to a pool using their Windows

domain credentials.

Note:

You can enable LDAP channel binding and LDAP signing on your AD domain controllers. For more information, see [Microsoft Security Advisory ADV190023](#).

You can configure varying levels of access for specific users by enabling Active Directory authentication, adding user accounts, and assigning roles to those accounts.

Active Directory users can use the xe CLI (passing appropriate `-u` and `-pw` arguments) and also connect to the host using XenCenter. Authentication is done on a per-resource pool basis.

Subjects control access to user accounts. A subject in XenServer maps to an entity on your Active Directory server (either a user or a group). When you enable external authentication, XenServer checks the credentials used to create a session against the local root credentials and then against the subject list. To permit access, create a subject entry for the person or group you want to grant access to. You can use XenCenter or the xe CLI to create a subject entry.

If you are familiar with XenCenter, note that the xe CLI uses slightly different terminology to refer to Active Directory and user account features:

XenCenter Term	xe CLI Term
Users, Add users	Subjects, Add subjects

Even though XenServer is Linux-based, XenServer lets you use Active Directory accounts for XenServer user accounts. To do so, it passes Active Directory credentials to the Active Directory domain controller.

When you add Active Directory to XenServer, Active Directory users and groups become XenServer subjects. The subjects are referred to as users in XenCenter. Users/groups are authenticated by using Active Directory on logon when you register a subject with XenServer. Users and groups do not need to qualify their user name by using a domain name.

To log on to a XenServer host, Active Directory users must have permission at the domain level to log on to the computer that hosts XenServer's machine account. By default, in a Windows Server 2019 domain, all users are permitted to log on to any computer in the domain. However, if you have changed this setting, ensure that the users that you want to have access to a XenServer host are permitted to log on at the domain level.

To qualify a user name, you must type the user name in Down-Level log on Name format, for example, `mydomain\myuser`.

Note:

By default, if you did not qualify the user name, XenCenter attempts to log in users to AD authentication servers using the domain to which it is joined. The exception to this is the LSU account, which XenCenter always authenticates locally (that is, on the XenServer) first.

The external authentication process works as follows:

1. The credentials supplied when connecting to a host are passed to the Active Directory domain controller for authentication.
2. The domain controller checks the credentials. If they are invalid, the authentication fails immediately.
3. If the credentials are valid, the Active Directory controller is queried to get the subject identifier and group membership associated with the credentials.
4. If the subject identifier matches the one stored in the XenServer, authentication succeeds.

When you join a domain, you enable Active Directory authentication for the pool. However, when a pool joins a domain, only users in that domain (or a domain with which it has trust relationships) can connect to the pool.

Note:

Manually updating the DNS configuration of a DHCP-configured network PIF is unsupported and can cause AD integration, and therefore user authentication, to fail or stop working.

Configure Active Directory authentication

XenServer supports use of Active Directory servers using Windows Server 2016 or later.

Note:

XenServer 9 does not support legacy AD ciphersuites such as RC4_HMAC_MD5, DES_CBC_MD5 and DES_CBC_CRC. You should ensure that your AD server supports at least one of AES128_HMAC_SHA1 or AES256_HMAC_SHA1

To authenticate Active Directory for XenServer hosts, you must use the same DNS server for both the Active Directory server (configured to allow interoperability) and the XenServer host.

In some configurations, the Active Directory server can provide the DNS itself. This can be achieved either using DHCP to provide the IP address and a list of DNS servers to the XenServer host. Alternatively, you can set the values in the PIF objects or use the installer when a manual static configuration is used.

We recommend enabling DHCP to assign host names. Do not assign the hostnames `localhost` or `linux` to hosts.

Warning:

XenServer host names must be unique throughout the XenServer deployment.

Note the following:

- XenServer labels its AD entry on the AD database using its hostname. If two XenServer hosts with the same hostname are joined to the same AD domain, the second XenServer overwrites the AD entry of the first XenServer. The overwriting occurs regardless of whether the hosts belong to the same or different pools. This can cause the AD authentication on the first XenServer to stop working.

You can use the same host name in two XenServer hosts, as long as they join different AD domains.

- The XenServer hosts can be in different time-zones, because it is the UTC time that is compared. To ensure that synchronization is correct, you can use the same NTP servers for your XenServer pool and the Active Directory server.
- Mixed-authentication pools are not supported. You cannot have a pool where some hosts in the pool are configured to use Active Directory and some are not.
- The XenServer Active Directory integration uses the Kerberos protocol to communicate with the Active Directory servers. Therefore, XenServer does not support communicating with Active Directory servers that do not use Kerberos.
- For external authentication using Active Directory to be successful, clocks on your XenServer hosts must be synchronized with the clocks on your Active Directory server. When XenServer joins the Active Directory domain, the synchronization is checked and authentication fails if there is too much skew between the servers.

Warning:

Host names must not be purely numeric. Because XenServer registers in Active Directory using the host's NetBIOS name, the host name should be no more than 15 characters (alphanumeric characters and hyphens, not starting or ending with a hyphen).

A limitation in recent SSH clients means that SSH does not work for usernames that contain any of the following characters: { } [] | &. Ensure that your usernames and Active Directory server names do not contain any of these characters.

When you add a host to a pool after enabling Active Directory authentication, you are prompted to configure Active Directory on the host joining the pool. When prompted for credentials on the joining host, type Active Directory credentials with sufficient privileges to add hosts to that domain.

Active Directory integration

Ensure that the following ports are open for connections outbound from XenServer to the Active Directory domain controllers.

Port	Protocol	Use
53	UDP/TCP	DNS
88	UDP/TCP	Kerberos 5
123	UDP	NTP
135	TCP	RPC Endpoint Mapper
137	UDP	NetBIOS Name Service
139	TCP	NetBIOS Session (SMB)
389	UDP/TCP	LDAP
445	TCP	SMB over TCP
464	UDP/TCP	Machine password changes
636	UDP/TCP	LDAP over SSL
3268	TCP	Global Catalog Search
49152-65535	TCP	RPC dynamic connections

For more information, see [Communication Ports Used by XenServer](#).

Note:

If users from a trusted domain need to log in to XenServer, ensure that the preceding ports are also open for connections outbound from XenServer to the domain controllers of the trusted domain.

Winbind

XenServer uses Winbind for authenticating Active Directory (AD) users with the AD server and to encrypt communications with the AD server.

Winbind does not support the following scenarios:

- Space at the beginning or end of a domain user or domain group name.
- Domain user names that contain 64 characters or more.
- Domain user names that include any of the special characters +<>`=/%@:;,\`
- Domain group names that include any of the special characters ,;,\`

Configuring Winbind Configure Winbind behavior with the following configuration options, which can be included in the `/etc/xapi.conf` file:

- `winbind_machine_pwd_timeout`: The value of this option defines how often, in seconds, the machine password is rotated for this XenServer host. Define a value as an integer.

The default value is 1209600 seconds (14 days). We recommend that you keep the default value or do not decrease the value below the default value to guarantee enough time to synchronize the new password between domain controllers.

- `winbind_kerberos_encryption_type`: The values for this option are `strong`, `legacy`, and `all`. The default value is `strong`.

- The value `all` permits the following cipher suites: `aes256-cts-hmac-sha1-96`, `aes128-cts-hmac-sha1-96`, and `arcfour-hmac-md5`

- The value `strong` permits the following cipher suites: `aes256-cts-hmac-sha1-96` and `aes128-cts-hmac-sha1-96`

- The value `legacy` permits the following cipher suites: `arcfour-hmac-md5`

The legacy option is insecure and we recommend that you only use it to debug issues.

For improved security, we recommend that you enforce AES encryption. To do this,

1. Ensure the domain controller supports `aes256-cts-hmac-sha1-96` and `aes128-cts-hmacsha1-96`.
2. Configure the domain controller to enable **The other domain supports Kerberos AES Encryption** in domain trust.

For more information, see [Method 3: Configure the trust to support AES128 and AES 256 encryption instead of RC4 encryption in the Microsoft documentation](#).

3. Update the `winbind_kerberos_encryption_type` option to use the value `strong`.
4. Restart the toolstack.

Do not restart the toolstack while HA is enabled. If possible, temporarily disable HA before restarting the toolstack.

- `winbind_set_machine_account_kerberos_encryption_type`: The values for this option are `true` and `false`. The default value is `false`.
 - The value `true` sets `msDS-SupportedEncryptionTypes` to `strong` on the Active Directory computer object for the XenServer host.
 - The value `false` does not configure `msDS-SupportedEncryptionTypes` on the Active Directory computer object for the XenServer host.

- `winbind_cache_time`: Winbind caches some domain information locally. The value of this option defines the number of seconds between each cache refresh. The default is 60 seconds.

After you update any of these configuration options, restart the toolstack.

How does XenServer manage the machine account password for AD integration?

Similarly to Windows client machines, Winbind automatically updates the machine account password. Winbind automatically updates the machine account password every 14 days or as specified by the configuration option `winbind_machine_pwd_timeout`.

Enable external authentication on a pool

External authentication using Active Directory can be configured using either XenCenter or the CLI using the following command.

```
1 xe pool-enable-external-auth auth-type=AD \  
2   service-name=full-qualified-domain \  
3   config:user=username \  
4   config:pass=password
```

The user specified must have `Add/remove computer objects or workstations` privilege, which is the default for domain administrators.

If you are not using DHCP on the network used by Active Directory and your XenServer hosts, use the following approaches to set up your DNS:

1. Set up your domain DNS suffix search order for resolving non-FQDN entries:

```
1 xe pif-param-set uuid=pif_uuid_in_the_dns_subnetwork \  
2   "other-config:domain=suffix1.com suffix2.com suffix3.com"
```

2. Configure the DNS server to use on your XenServer hosts:

```
1 xe pif-reconfigure-ip mode=static dns=dnshost ip=ip \  
2   gateway=gateway netmask=netmask uuid=uuid
```

3. Manually set the management interface to use a PIF that is on the same network as your DNS server:

```
1 xe host-management-reconfigure pif-uuid=pif_in_the_dns_subnetwork
```

Note:

External authentication is a per-host property. However, we recommend that you enable and disable external authentication on a per-pool basis. A per-pool setting allows XenServer to deal

with failures that occur when enabling authentication on a particular host. XenServer also rolls back any changes that may be required, ensuring a consistent configuration across the pool. Use the `host-param-list` command to inspect properties of a host and to determine the status of external authentication by checking the values of the relevant fields.

Use XenCenter to disable Active Directory authentication, or the following `xe` command:

```
1 xe pool-disable-external-auth
```

Enable AD authentication caching

The Windows AD infrastructure contains intrinsic delays when replicating information about users, such as blocked/unblocked, password, and other fields during replication across different AD sites. This can lead to AD authentication taking a very long time, especially when operating at scale in global AD deployments. Enabling authentication caching allows the system to remember authentication decisions for a limited time, helping to speed up login when external Active Directory (AD) authentication is slow. This feature is disabled by default.

To enable AD authentication caching:

```
1 xe pool-param-set uuid=<pool-uuid> ext-auth-cache-enabled=true
```

By default, authentication decisions are remembered for 300 seconds. This can be adjusted:

```
1 xe pool-param-set uuid=<pool-uuid> ext-auth-cache-expiry=<seconds>
```

To disable AD authentication caching:

```
1 xe pool-param-set uuid=<pool-uuid> ext-auth-cache-enabled=false
```

User authentication

To allow a user access to your XenServer host, you must add a subject for that user or a group that they are in. (Transitive group memberships are also checked in the normal way. For example, adding a subject for group **A**, where group **A** contains group **B** and **user 1** is a member of group **B** would permit access to **user 1**.) If you want to manage user permissions in Active Directory, you can create a single group that you then add and delete users to/from. Alternatively, you can add and delete individual users from XenServer, or a combination of users and groups as appropriate for your authentication requirements. You can manage the subject list from XenCenter or using the CLI as described in the following section.

When authenticating a user, the credentials are first checked against the local root account, allowing you to recover a system whose AD server has failed. If the credentials (user name and password) do

not match, then an authentication request is made to the AD server. If the authentication is successful, the user's information is retrieved and validated against the local subject list. Access is denied if the authentication fails. Validation against the subject list succeeds if the user or a group in the transitive group membership of the user is in the subject list.

Note:

When using Active Directory groups to grant access for Pool Administrator users who require host ssh access, the size of the AD group must not exceed 500 users.

To add an AD subject to XenServer:

```
1 xe subject-add subject-name=entity_name
```

The `entity_name` is the name of the user or group to which you want to grant access. You can include the domain of the entity (for example, 'xendrt\user1' as opposed to 'user1') although the behavior is the same unless disambiguation is required.

Find the user's subject identifier. The identifier is the user or the group containing the user. Removing a group removes access to all users in that group, provided they are not also specified in the subject list. Use the `subject list` command to find the user's subject identifier. :

```
1 xe subject-list
```

This command returns a list of all users.

To apply a filter to the list, for example to find the subject identifier for a user `user1` in the `testad` domain, use the following command:

```
1 xe subject-list other-config:subject-name='testad\user1'
```

Remove the user using the `subject-remove` command, passing in the subject identifier you learned in the previous step:

```
1 xe subject-remove subject-uuid=subject_uuid
```

You can end any current session this user has already authenticated. For more information, see *Terminating all authenticated sessions using xe* and *Terminating individual user sessions using xe* in the following section. If you do not end sessions, users with revoked permissions may continue to access the system until they log out.

Run the following command to identify the list of users and groups with permission to access your XenServer host or pool:

```
1 xe subject-list
```

Remove access for a user

When a user is authenticated, they can access the host until they end their session, or another user ends their session. Removing a user from the subject list, or removing them from a group in the subject list, doesn't automatically revoke any already-authenticated sessions that the user has. Users can continue to access the pool using XenCenter or other API sessions that they have already created. XenCenter and the CLI provide facilities to end individual sessions, or all active sessions forcefully. See the [XenCenter documentation](#) for information on procedures using XenCenter, or the following section for procedures using the CLI.

Terminate all authenticated sessions using xe

Run the following CLI command to end all authenticated sessions using xe:

```
1 xe session-subject-identifier-logout-all
```

Terminate individual user sessions using xe

1. Determine the subject identifier whose session you want to log out. Use either the `session-subject-identifier-list` or `subject-list` xe commands to find the subject identifier. The first command shows users who have sessions. The second command shows all users but can be filtered. For example, by using a command like `xe subject-list other-config:subject-name=xendt\\user1`. You may need a double backslash as shown depending on your shell).
2. Use the `session-subject-logout` command, passing the subject identifier you have determined in the previous step as a parameter, for example:

```
1 xe session-subject-identifier-logout subject-identifier=subject_id
```

Leave an AD domain

Warning:

When you leave the domain, any users who authenticated to the pool or host with Active Directory credentials are disconnected.

Use XenCenter to leave an AD domain. For more information, see the [XenCenter documentation](#). Alternately run the `pool-disable-external-auth` command, specifying the pool UUID if necessary.

Note:

Leaving the domain does not delete the host objects from the AD database. Refer to the Active Directory documentation for information about how to detect and remove your disabled host entries.

Role-based access control

December 17, 2025

The Role-Based Access Control (RBAC) feature in XenServer® allows you to assign users, roles, and permissions to control who has access to your XenServer and what actions they can perform. The XenServer RBAC system maps a user (or a group of users) to defined roles (a named set of permissions). The roles have associated XenServer permissions to perform certain operations.

Permissions are not assigned to users directly. Users acquire permissions through roles assigned to them. Therefore, managing individual user permissions becomes a matter of assigning the user to the appropriate role, which simplifies common operations. XenServer maintains a list of authorized users and their roles.

RBAC allows you to restrict which operations different groups of users can perform, reducing the probability of an accident by an inexperienced user.

RBAC also provides an Audit Log feature for compliance and auditing.

RBAC depends on Active Directory for authentication services. Specifically, XenServer keeps a list of authorized users based on Active Directory user and group accounts. As a result, you must join the pool to the domain and add Active Directory accounts before you can assign roles.

The local super user (LSU), or root, is a special user account used for system administration and has all rights or permissions. The local super user is the default account at installation in XenServer. The LSU is authenticated through XenServer and not through an external authentication service. If the external authentication service fails, the LSU can still log in and manage the system. The LSU can always access the XenServer physical host through SSH.

RBAC process

The following section describes the standard process for implementing RBAC and assigning a user or group a role:

1. Join the domain. For more information, see [Enabling external authentication on a pool](#).

2. Add an Active Directory user or group to the pool. This becomes a subject. For more information, see [To add a subject to RBAC](#).
3. Assign (or change) the subject's RBAC role. For more information, see [To assign an RBAC role to a subject](#).

RBAC roles and permissions

March 16, 2026

Roles

XenServer® is shipped with the following six, pre-established roles:

- *Pool Administrator* (Pool Admin) –the same as the local root. Can perform all operations.

Note:

The local super user (root) has the “Pool Admin” role. The Pool Admin role has the same permissions as the local root.

If you remove the Pool Admin role from a user, consider also changing the root password and rotating the pool secret. For more information, see [Pool Security](#).

- *Pool Operator* (Pool Operator) –can do everything apart from adding/removing users and changing their roles. This role is focused mainly on host and pool management (that is, creating storage, making pools, managing the hosts and so on.)
- *Virtual Machine Power Administrator* (VM Power Admin) –creates and manages Virtual Machines. This role is focused on provisioning VMs for use by a VM operator.
- *Virtual Machine Administrator* (VM Admin) –similar to a VM Power Admin, but cannot migrate VMs or perform snapshots.
- *Virtual Machine Operator* (VM Operator) –similar to VM Admin, but cannot create/destroy VMs – but can perform start/stop lifecycle operations.
- *Read-only* (Read Only) –can view resource pool and performance data.

Note:

- To apply updates to XenServer 9 pools, you must be logged in to XenCenter as a Pool Administrator or Pool Operator, or using a local root account.

- When using Active Directory groups to grant access for Pool Administrator users who require host SSH access, the number of users in the Active Directory group must not exceed 500.

For a summary of the permissions available for each role and for information on the operations available for each permission, see *Definitions of RBAC roles and permissions* in the following section.

When you create a user in XenServer, you must first assign a role to the newly created user before they can use the account. XenServer **does not** automatically assign a role to the newly created user. As a result, these accounts do not have any access to XenServer pool until you assign them a role.

1. Modify the subject to role mapping. This requires the assign/modify role permission, only available to a Pool Administrator.
2. Modify the user’s containing group membership in Active Directory.

Definitions of RBAC roles and permissions

The following table summarizes which permissions are available for each role. For details on the operations available for each permission, see *Definitions of permissions*.

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Assign/modifyX roles						
Log in to (physical) server consoles (through SSH and XenCenter)	X					
Server backup/re- store	X					
Install a TLS certificate on a server	X					
Apply updates to a pool	X	X				

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Rolling Pool Upgrade	X					
Import OVF/OVA packages	X	X				
Import XVA packages; import disk images	X	X	X			
Export OVF/O-VA/XVA packages	X	X	X	X		
Set cores per socket	X	X	X	X		
Convert VMs using Conversion Manager	X					
Switch-port locking	X	X				
Multipathing	X	X				
Log out active user connections	X	X				
Monitor host and dom0 resources with NRPE	X					
Monitor host and dom0 resources with SNMP	X					

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Create and dismiss alerts	X	X				
Cancel task of any user	X	X				
Pool man- agement	X	X				
Live migration	X	X	X			
Storage live migration	X	X	X			
VM advanced operations	X	X	X			
VM cre- ate/destroy operations	X	X	X	X		
VM change CD media	X	X	X	X	X	
VM change power state	X	X	X	X	X	
View VM consoles	X	X	X	X	X	
VM groups (Add/Modi- fy/Delete VM groups)	X	X	X	X		
VM groups (Add/Re- move VMs to an existing VM group)	X	X	X	X		
VM groups (View VM groups)	X	X	X	X	X	X

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
vApps (Add/Modi- fy/Delete vApps)	X	X				
vApps (Start/Shut- down vApps)	X	X				
vApps (Ad- d/Remove VMs to an existing vApp)	X	X				
vApps (View vApps)	X	X	X	X	X	X
XenCenter view man- agement operations	X	X	X	X	X	
Cancel own tasks	X	X	X	X	X	X
Read audit logs	X	X	X	X	X	X
Configure, initialize, enable, disable Workload Balancing (WLB)	X	X				
Apply WLB optimiza- tion recommen- dations	X	X				

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Accept WLB placement recommen- dations	X	X	X			
Display WLB config- uration	X	X	X	X	X	X
Generate WLB reports	X	X	X	X	X	X
Connect to pool and read all pool metadata (including viewing alerts)	X	X	X	X	X	X
Configure virtual GPU	X	X				
View virtual GPU config- uration	X	X	X	X	X	X
Gather diagnostic information	X	X				
vCPU Hotplug	X	X	X	X		
Configure changed block tracking	X	X	X	X		
List changed blocks	X	X	X	X	X	

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Configure PVS-Accelerator	X	X				
View PVS-Accelerator configuration	X	X	X	X	X	X
Scheduled Snapshots (Add/Remove VMs to existing Snapshots Schedules)	X	X	X			
Scheduled Snapshots (Add/Modify/Delete Snapshot Schedules)	X	X				

Definitions of permissions

Assign/modify roles:

- Add and remove users
- Add and remove roles from users
- Enable and disable Active Directory integration (being joined to the domain)

This permission lets the user grant themselves any permission or perform any task.

Warning:

This role lets the user disable the Active Directory integration and all subjects added from Active Directory.

Log in to server consoles:

- Server console access through ssh

- Server console access through XenCenter®

Warning:

With access to a root shell, the assignee can arbitrarily reconfigure the entire system, including RBAC.

Server backup/restore:

- Back up and restore servers
- Back up and restore pool metadata

The ability to restore a backup lets the assignee revert RBAC configuration changes.

Install a TLS certificate on a server:

This permission enables an administrator to install a TLS certificate on a server.

Apply updates to a pool:

- Synchronize your pool with the content delivery network (CDN)
- Apply the updates by migrating VMs off each host if necessary and running any necessary update tasks such as rebooting the host, restarting the toolstack, or rebooting the VMs

Rolling Pool Upgrade:

- Upgrade all hosts in a pool using the **Rolling Pool Upgrade** wizard.

Import OVF/OVA packages:

- Import OVF and OVA packages

Import XVA packages and disk images:

- Import XVA packages
- Import disk images

Export OVF/OVA/XVA packages:

- Export VMs as OVF/OVA packages
- Export VMs as XVA packages

Set cores-per-socket:

- Set the number of cores per socket for the VM's virtual CPUs

This permission enables the user to specify the topology for the VM's virtual CPUs.

Convert VMs using Conversion Manager:

- Convert VMware ESXi/vCenter VMs to XenServer VMs

This permission lets the user convert workloads from VMware to XenServer. Convert these workloads by copying batches of VMware ESXi/vCenter VMs to the XenServer environment.

Switch-port locking:

- Control traffic on a network

This permission lets the user block all traffic on a network by default, or define specific IP addresses from which a VM can send traffic.

Multipathing:

- Enable multipathing
- Disable multipathing

Log out active user connections:

- Ability to disconnect logged in users

Monitor host and dom0 resources with NRPE:

For more information, see [Monitor host and dom0 resources with NRPE](#).

Monitor host and dom0 resources with SNMP:

For more information, see [Monitor host and dom0 resources with SNMP](#).

Create/dismiss alerts:

- Configure XenCenter to generate alerts when resource usage crosses certain thresholds
- Remove alerts from the Alerts view

Warning: A user with this permission can dismiss alerts for the entire pool.

Note:

The ability to view alerts is part of the **Connect to pool and read all pool metadata permission**.

Cancel task of any user:

- Cancel any user's running task

This permission lets the user request XenServer cancel an in-progress task initiated by any user.

Pool management:

- Set pool properties (naming, default SRs)
- Create a GFS2 clustered pool
- Enable, disable, and configure HA

- Set per-VM HA restart priorities
- Configure DR and perform DR failover, failback, and test failover operations.
- Enable, disable, and configure Workload Balancing (WLB)
- Add and remove server from pool
- Emergency transition to pool coordinator
- Emergency pool coordinator address
- Emergency recovery of pool members
- Designate new pool coordinator
- Manage pool and server certificates
- Patching
- Set server properties
- Configure server logging
- Enable and disable servers
- Shut down, reboot, and power-on servers
- Restart toolstack
- System status reports
- Apply license
- Live migration of all other VMs on a server to another server, due to either WLB, maintenance mode, or high availability
- Configure server management interfaces and secondary interfaces
- Disable server management
- Delete crashdumps
- Add, edit, and remove networks
- Add, edit, and remove PBDs/PIFs/VLANs/Bonds/SRs
- Add, remove, and retrieve secrets

This permission includes all the actions required to maintain a pool.

Note: If the management interface is not functioning, no logins can authenticate except local root logins.

Live migration:

- Migrate VMs from one host to another host when the VMs are on storage shared by both hosts

Storage live migration:

- Migrate from one host to another host when the VMs are not on storage shared between the two hosts
- Move Virtual Disk (VDIs) from one SR to another SR

VM advanced operations:

- Adjust VM memory (through Dynamic Memory Control)

- Create a VM snapshot with memory, take VM snapshots, and roll-back VMs
- Migrate VMs
- Start VMs, including specifying physical server
- Resume VMs

This permission provides the assignee with enough privileges to start a VM on a different host if they are not satisfied with the host XenServer selected.

VM create/destroy operations:

- Install and delete VMs
- Clone/copy VMs
- Add, remove, and configure virtual disk/CD devices
- Add, remove, and configure virtual network devices
- VM configuration change
- VM metadata backup/restore
- Convert a VM into a template
- Convert an existing snapshot into a template

VM change CD media:

- Eject current CD
- Insert new CD

VM change power state:

- Start VMs (automatic placement)
- Shut down VMs
- Reboot VMs
- Suspend VMs
- Resume VMs (automatic placement)

This permission does not include `start_on`, `resume_on`, and `migrate`, which are part of the VM advanced operations permission.

View VM consoles:

- See and interact with VM consoles

This permission does not let the user view host consoles.

VM groups (Add/Modify/Delete groups):

- Create a VM group
- Delete a VM group
- Change the properties of a VM group

VM groups (Add/Remove VMs to an existing VM group):

- Assign a VM to a VM group
- Remove a VM from a VM group
- Change the VM group a VM is assigned to

VM groups (View VM groups):

- View VM groups in the pool

vApps (Add/Modify/Delete vApps):

- Create a vApp
- Delete a vApp
- Change the properties of a vApp

vApps (Start/Shutdown vApps):

- Start a vApp
- Shutdown a vApp

vApps (Add/Remove VMs to an existing vApp):

- Add a VM to a vApp.
- Remove a VM from a vApp

vApps (View vApps):

- View vApps in the pool

Cancel own tasks:

- Enables users to cancel their own tasks

Read audit log:

- Download XenServer audit log

Configure, initialize, enable, disable WLB:

- Configure WLB
- Initialize WLB and change WLB servers
- Enable WLB
- Disable WLB

Apply WLB optimization recommendations:

- Apply any optimization recommendations that appear in the **WLB** tab

Modify WLB report subscriptions:

- Change the WLB report generated or its recipient

Accept WLB placement recommendations:

- Select one of the servers Workload Balancing recommends for placement (“star” recommendations)

Display WLB configuration:

- View WLB settings for a pool as shown on the **WLB** tab

Generate WLB reports:

- View and run WLB reports, including the Pool Audit Trail report

XenCenter view management operations:

- Create and modify global XenCenter folders
- Create and modify global XenCenter custom fields
- Create and modify global XenCenter searches

Connect to pool and read all pool metadata:

- Log in to pool
- View pool metadata
- View historical performance data
- View logged in users
- View users and roles
- View tasks
- View messages
- Register for and receive events

Configure virtual GPU:

- Specify a pool-wide placement policy
- Assign a virtual GPU to a VM
- Remove a virtual GPU from a VM
- Modify allowed virtual GPU types
- Create, destroy, or assign a GPU group

View virtual GPU configuration:

- View GPUs, GPU placement policies, and virtual GPU assignments.

Gather diagnostic information from XenServer:

- Initiate GC collection and heap compaction
- Gather garbage collection statistics
- Gather database statistics
- Gather network statistics

Configure changed block tracking:

- Enable changed block tracking
- Disable changed block tracking
- Destroy the data associated with a snapshot and retain the metadata
- Get the NBD connection information for a VDI
- Export a VDI over an NBD connection

Changed block tracking can be enabled only for licensed instances of XenServer Premium Edition.

List changed blocks:

- Compare two VDI snapshots and list the blocks that have changed between them.

Configure PVS-Accelerator:

- Enable PVS-Accelerator
- Disable PVS-Accelerator
- Update PVS-Accelerator cache configuration
- Add or Remove PVS-Accelerator cache configuration

View PVS-Accelerator configuration:

- View the status of PVS-Accelerator

Scheduled snapshots (Add/Remove VMs to existing Snapshots Schedules):

- Add VMs to existing snapshot schedules
- Remove VMs from existing snapshot schedules

Scheduled snapshots (Add/Modify/Delete Snapshot Schedules):


- Add snapshot schedules
- Modify snapshot schedules
- Delete snapshot schedules

Note:

Sometimes, a Read Only user cannot move a resource into a folder in XenCenter, even after receiving an elevation prompt and supplying the credentials of a more privileged user. In this case, log on to XenCenter as the more privileged user and retry the action.

How does XenServer compute the roles for the session?

1. The subject is authenticated through the Active Directory server to verify which containing groups the subject may also belong to.
2. XenServer then verifies which roles have been assigned both to the subject, and to its containing groups.
3. As subjects can be members of multiple Active Directory groups, they inherit all of the permissions of the associated roles.



Use RBAC with the CLI

December 17, 2025

RBAC xe CLI commands

Use the following commands to work with roles and subjects.

To list all the available defined roles

Run the command: `xe role-list`

This command returns a list of the currently defined roles, for example:

```
1      uuid( R0): 0165f154-ba3e-034e-6b27-5d271af109ba
2      name ( R0): pool-admin
3      description ( R0): The Pool Administrator role has full access to
4          all
5          features and settings, including accessing Dom0 and managing
6          subjects,
7          roles and external authentication
8
9      uuid ( R0): b9ce9791-0604-50cd-0649-09b3284c7dfd
10     name ( R0): pool-operator
11     description ( R0): The Pool Operator role manages host- and pool-
12         wide resources,
13         including setting up storage, creating resource pools and managing
14         patches, and
15         high availability (HA).
16
17     uuid( R0): 7955168d-7bec-10ed-105f-c6a7e6e63249
18     name ( R0): vm-power-admin
19     description ( R0): The VM Power Administrator role has full access
20         to VM and
21         template management and can choose where to start VMs and use the
22         dynamic memory
23         control and VM snapshot features
24
25     uuid ( R0): aaa00ab5-7340-bfbc-0d1b-7cf342639a6e
26     name ( R0): vm-admin
27     description ( R0): The VM Administrator role can manage VMs and
28         templates
29
30     uuid ( R0): fb8d4ff9-310c-a959-0613-54101535d3d5
31     name ( R0): vm-operator
32     description ( R0): The VM Operator role can use VMs and interact
33         with VM consoles
34
35     uuid ( R0): 7233b8e3-eacb-d7da-2c95-f2e581cdbf4e
36     name ( R0): read-only
37     description ( R0): The Read-Only role can log in with basic read-
38         only access
```

Note:

This list of roles is static. You cannot add, remove, or modify roles.

To display a list of current subjects

Run the following command:

```
1 xe subject-list
```

This command returns a list of XenServer® users, their uuid, and the roles they are associated with:

```
1      uuid ( RO): bb6dd239-1fa9-a06b-a497-3be28b8dca44
2      subject-identifier ( RO): S
      -1-5-21-1539997073-1618981536-2562117463-2244
3      other-config (MRO): subject-name: example01\user_vm_admin; subject-
      upn: \
4      user_vm_admin@XENDT.NET; subject-uid: 1823475908; subject-gid:
      1823474177; \
5      subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2244;
      subject-gecos: \
6      user_vm_admin; subject-displayname: user_vm_admin; subject-is-
      group: false; \
7      subject-account-disabled: false; subject-account-expired: false;
      \
8      subject-account-locked: false; subject-password-expired: false
9      roles (SRO): vm-admin
10
11     uuid ( RO): 4fe89a50-6a1a-d9dd-afb9-b554cd00c01a
12     subject-identifier ( RO): S
      -1-5-21-1539997073-1618981536-2562117463-2245
13     other-config (MRO): subject-name: example02\user_vm_op; subject-upn
      : \
14     user_vm_op@XENDT.NET; subject-uid: 1823475909; subject-gid:
      1823474177; \
15     subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2245; \
16     subject-gecos: user_vm_op; subject-displayname: user_vm_op; \
17     subject-is-group: false; subject-account-disabled: false; \
18     subject-account-expired: false; subject-account-locked: \
19     false; subject-password-expired: false
20     roles (SRO): vm-operator
21
22     uuid ( RO): 8a63fbf0-9ef4-4fef-b4a5-b42984c27267
23     subject-identifier ( RO): S
      -1-5-21-1539997073-1618981536-2562117463-2242
24     other-config (MRO): subject-name: example03\user_pool_op; \
25     subject-upn: user_pool_op@XENDT.NET; subject-uid: 1823475906; \
26     subject-gid: 1823474177; subject-s id:
27     S-1-5-21-1539997073-1618981536-2562117463-2242; \
28     subject-gecos: user_pool_op; subject-displayname: user_pool_op; \
29     subject-is-group: false; subject-account-disabled: false; \
30     subject-account-expired: false; subject-account-locked: \
31     false; subject-password-expired: false
32     roles (SRO): pool-operator
```

To add a subject to RBAC

To enable existing AD users to use RBAC, create a subject instance within XenServer, either for the AD user directly, or for the containing groups:

Run the following command to add a new subject instance:

```
1 xe subject-add subject-name=AD user/group
```

To assign an RBAC role to a subject

After adding a subject, you can assign it to an RBAC role. You can refer to the role by either by its UUID or name:

Run the command:

```
1 xe subject-role-add uuid=subject uuid role-uuid=role_uuid
```

Or

```
1 xe subject-role-add uuid=subject uuid role-name=role_name
```

For example, the following command adds a subject with the UUID `b9b3d03b-3d10-79d3-8ed7-a782c5ea13b4` to the Pool Administrator role:

```
1 xe subject-role-add uuid=b9b3d03b-3d10-79d3-8ed7-a782c5ea13b4 role-name=pool-admin
```

To change the RBAC role of a subject

To change the role of a user, it is necessary to remove them from their existing role and add them to a new role:

Run the following commands:

```
1 xe subject-role-remove uuid=subject_uuid role-name=role_name_to_remove
2 xe subject-role-add uuid=subject_uuid role-name=role_name_to_add
```

The user must log out and log back in to ensure that the new role takes effect. This requires the “Logout Active User Connections” permission available to a Pool Administrator or Pool Operator.

If you remove the Pool Admin role from a user, consider also changing the root password and rotating the pool secret. For more information, see [Pool Security](#).

Warning:

When you add or remove a pool-admin subject, it can take a few seconds for all hosts in the pool

to accept ssh sessions associated with this subject.

Auditing

The RBAC audit log records any operation taken by a logged-in user.

- The message records the Subject ID and user name associated with the session that invoked the operation.
- If a subject invokes an operation that is not authorized, the operation is logged.
- Any successful operation is also recorded. If the operation failed then the error code is logged.

Audit log xe CLI commands

The following command downloads all the available records of the RBAC audit file in the pool to a file. If the optional parameter 'since' is present, then it only downloads the records from that specific point in time.

```
1 xe audit-log-get \[since=timestamp\] filename=output filename
```

To obtain all audit records from the pool

Run the following command:

```
1 xe audit-log-get filename=/tmp/auditlog-pool-actions.out
```

To obtain audit records of the pool since a precise millisecond timestamp

Run the following command:

```
1 xe audit-log-get since=2009-09-24T17:56:20.530Z \  
2     filename=/tmp/auditlog-pool-actions.out
```

To obtain audit records of the pool since a precise minute timestamp

Run the following command:

```
1 xe audit-log-get since=2009-09-24T17:56Z \  
2     filename=/tmp/auditlog-pool-actions.out
```


Secure Boot for XenServer 9

June 18, 2026

Secure Boot helps prevent untrusted code from running during host boot and runtime by enforcing a verified chain of trust. XenServer 9 supports installation and operation with Secure Boot enabled in firmware.

Benefits

- Enforces signature verification for each boot component.
- Reduces risk of executing untrusted or malicious code.
- Improves overall platform integrity against low-level attacks.

Support and requirements

Supported scenarios in XenServer 9:

- Host installation from ISO or network
- Boot from SAN
- Normal host runtime

Requirements:

- UEFI firmware is required; hosts do not boot without UEFI.
- Enable Secure Boot in firmware settings.
- Certificate requirements:
 - Microsoft UEFI CA 2011 (or 2023, if available on your platform).
- All Dom0 kernel modules must be signed.

Notes:

- Host Secure Boot is independent of VM Secure Boot; enabling one does not enable the other.
- Mixed pools (hosts with Secure Boot both enabled and disabled) are supported.
- You can enable Secure Boot after an RPU from XenServer 8.4, during a manual upgrade from XenServer 8.4, or after a fresh XenServer 9 install.
- You can disable Secure Boot at any time in firmware.
- If you restore a host from XenServer 9 back to XenServer 8.4, disable Secure Boot before booting.
- Running memory tests from the boot menu requires Secure Boot to be disabled.
- Some unsupported Xen or Dom0 kernel command-line options might not work with Secure Boot enabled.

Configure Secure Boot

Configuration steps vary by platform. Refer to your hardware vendor documentation for exact steps. The following example uses a Supermicro server.

1. In firmware setup, go to Security > Secure Boot.
2. Set CSM Support to Disabled.
3. Set Secure Boot Mode to Standard.
4. Set Secure Boot Control to Enabled.

Networking

June 22, 2026

This section provides an overview of XenServer® networking, including networks, VLANs, and NIC bonds. It also discusses how to manage your networking configuration and troubleshoot it.

If you are already familiar with XenServer networking concepts, you can skip ahead to [Manage networking](#) for information about the following sections:

- Create networks for standalone XenServer hosts
- Create networks for XenServer hosts that are configured in a resource pool
- Create VLANs for XenServer hosts, either standalone or part of a resource pool
- Create bonds for standalone XenServer hosts
- Create bonds for XenServer hosts that are configured in a resource pool

Note:

The term ‘management interface’ is used to indicate the IP-enabled NIC that carries the management traffic. The term ‘secondary interface’ is used to indicate an IP-enabled NIC configured for storage traffic.

Networking support

XenServer supports up to 16 physical network interfaces (or up to 4 bonded network interfaces) per host and up to 7 virtual network interfaces per VM.

Note:

XenServer provides automated configuration and management of NICs using the `xe` command

line interface (CLI). Do not edit the host networking configuration files directly.

XenServer networking overview

This section describes the general concepts of networking in the XenServer environment.

XenServer creates a network for each physical NIC during installation. When you add a host to a pool, the default networks are merged. This is to ensure all physical NICs with the same device name are attached to the same network.

Typically, you add a network to create an internal network, set up a new VLAN using an existing NIC, or create a NIC bond.

You can configure the following different types of networks in XenServer:

- **External networks** have an association with a physical network interface. External networks provide a bridge between a virtual machine and the physical network interface connected to the network. External networks enable a virtual machine to connect to resources available through the host's physical NIC.
- **Bonded networks** create a bond between two or more NICs to create a single, high-performing channel between the virtual machine and the network.
- **Single-Server Private networks** have no association to a physical network interface. Single-server private networks can be used to provide connectivity between the virtual machines on a given host, with no connection to the outside world.

Note:

Some networking options have different behaviors when used with standalone XenServer hosts compared to resource pools. This section contains sections on general information that applies to both standalone hosts and pools, followed by specific information and procedures for each.

Network objects

This section uses three types of server-side software objects to represent networking entities. These objects are:

- A *PIF*, which represents a physical NIC on a host. PIF objects have a name and description, a UUID, the parameters of the NIC they represent, and the network and host they are connected to.
- A *VIF*, which represents a virtual NIC on a virtual machine. VIF objects have a name and description, a UUID, and the network and VM they are connected to.

- A *network*, which is a virtual Ethernet switch on a host. Network objects have a name and description, a UUID, and the collection of VIFs and PIFs connected to them.

XenCenter® and the xe CLI allow you to configure networking options. You can control the NIC used for management operations, and create advanced networking features such as VLANs and NIC bonds.

Networks

Each XenServer host has one or more networks, which are virtual Ethernet switches. Networks that are not associated with a PIF are considered *internal*. Internal networks can be used to provide connectivity only between VMs on a given XenServer host, with no connection to the outside world. Networks associated with a PIF are considered *external*. External networks provide a bridge between VIFs and the PIF connected to the network, enabling connectivity to resources available through the PIF's NIC.

VLANs

VLANs, as defined by the IEEE 802.1Q standard, allow a single physical network to support multiple logical networks. XenServer hosts support VLANs in multiple ways.

Note:

- When using GFS2 SR and a clustered pool, the cluster network cannot be on a non-management VLAN.
- All supported VLAN configurations are equally applicable to pools and standalone hosts, and bonded and non-bonded configurations.

Using VLANs with virtual machines Switch ports configured as 802.1Q VLAN trunk ports can be used with the XenServer VLAN features to connect guest virtual network interfaces (VIFs) to specific VLANs. In this case, the XenServer host performs the VLAN tagging/untagging functions for the guest, which is unaware of any VLAN configuration.

XenServer VLANs are represented by additional PIF objects representing VLAN interfaces corresponding to a specified VLAN tag. You can connect XenServer networks to the PIF representing the physical NIC to see all traffic on the NIC. Alternatively, connect networks to a PIF representing a VLAN to see only the traffic with the specified VLAN tag. You can also connect a network such that it only sees the native VLAN traffic, by attaching it to VLAN 0.

For procedures on how to create VLANs for XenServer hosts, either standalone or part of a resource pool, see [Creating VLANs](#).

If you want the guest to perform the VLAN tagging and untagging functions, the guest must be aware of the VLANs. When configuring the network for your VMs, configure the switch ports as VLAN trunk ports, but do not create VLANs for the XenServer host. Instead, use VIFs on a normal, non-VLAN network.

Using VLANs with management interfaces Management interface can be configured on a VLAN using a switch port configured as trunk port or access mode port. Use XenCenter or xe CLI to set up a VLAN and make it the management interface. For more information, see [Management interface](#).

Using VLANs with dedicated storage NICs Dedicated storage NICs can be configured to use native VLAN or access mode ports as described in the previous section for management interfaces. Dedicated storage NICs are also known as IP-enabled NICs or secondary interfaces. You can configure dedicated storage NICs to use trunk ports and XenServer VLANs as described in the previous section for virtual machines. For more information, see [Configuring a dedicated storage NIC](#).

Combining management interfaces and guest VLANs on a single host NIC A single switch port can be configured with both trunk and native VLANs, allowing one host NIC to be used for a management interface (on the native VLAN) and for connecting guest VIFs to specific VLAN IDs.

Jumbo frames

Jumbo frames can be used to optimize the performance of traffic on storage networks, VM networks, and the management network. Jumbo frames are Ethernet frames containing more than 1,500 bytes of payload. Jumbo frames are typically used to achieve better throughput, reduce the load on system bus memory, and reduce the CPU overhead.

Requirements for using jumbo frames Note the following when using jumbo frames:

- Jumbo frames are configured at a pool level
- Every device on the subnet must be configured to use jumbo frames

To use jumbo frames, set the Maximum Transmission Unit (MTU) to a value between 1500 and 9216. You can use XenCenter or the xe CLI to set the MTU.

NIC Bonds

NIC bonds, sometimes also known as NIC teaming, improve XenServer host resiliency and bandwidth by enabling administrators to configure two or more NICs together. NIC bonds logically function as one network card and all bonded NICs share a MAC address.

If one NIC in the bond fails, the host's network traffic is automatically redirected through the second NIC. XenServer supports up to eight bonded networks.

XenServer supports active-active, active-passive, and LACP bonding modes. You can bond either two, three or four NICs.

All bonding modes support failover. However, not all modes allow all links to be active for all traffic types. XenServer supports bonding the following types of NICs together:

- **NICs (non-management).** You can bond NICs that XenServer is using solely for VM traffic. Bonding these NICs not only provides resiliency, but doing so also balances the traffic from multiple VMs between the NICs.
- **Management interfaces.** You can bond a management interface to another NIC so that the second NIC provides failover for management traffic. Although configuring an LACP link aggregation bond provides load balancing for management traffic, active-active NIC bonding does not. You can create a VLAN on bonded NICs and the host management interface can be assigned to that VLAN.
- **Secondary interfaces.** You can bond NICs that you have configured as secondary interfaces (for example, for storage). However, for most iSCSI software initiator storage, we recommend configuring multipathing instead of NIC bonding as described in the *Designing XenServer Network Configurations*.

Throughout this section, the term IP-based storage traffic is used to describe iSCSI and NFS traffic collectively.

You can create a bond if a VIF is already using one of the interfaces that will be bonded: the VM traffic migrates automatically to the new bonded interface.

In XenServer, An additional PIF represents a NIC bond. XenServer NIC bonds completely subsume the underlying physical devices (PIFs).

Notes:

- Creating a bond that contains only one NIC is not supported.
- The bonded NICs can be different models to each other.

Best practices

Follow these best practices when setting up your NIC bonds:

- Connect the links of the bond to different physical network switches, not just ports on the same switch.
- Ensure that the separate switches draw power from different, independent power distribution units (PDUs).

- If possible, in your data center, place the PDUs on different phases of the power feed or even feeds provided by different utility companies.
- Consider using uninterruptible power supply units to ensure that the network switches and hosts can continue to function or perform an orderly shutdown in the event of a power failure.

These measures add resiliency against software, hardware, or power failures that can affect your network switches.

Key points about IP addressing

Bonded NICs either have one IP address or no IP addresses, as follows:

- **Management and storage networks.**
 - If you bond a management interface or secondary interface, a single IP address is assigned to the bond. That is, each NIC does not have its own IP address. XenServer treats the two NICs as one logical connection.
 - When bonds are used for non-VM traffic, for example, to connect to shared network storage or XenCenter for management, configure an IP address for the bond. However, if you have already assigned an IP address to one of the NICs (that is, created a management interface or secondary interface), that IP address is assigned to the entire bond automatically.
 - If you bond a management interface or secondary interface to a NIC without an IP address, the bond assumes the IP address of the respective interface.
 - If you bond a tagged VLAN management interface and a secondary interface, the management VLAN is created on that bonded NIC.
- **VM networks.** When bonded NICs are used for VM traffic, you do not need to configure an IP address for the bond. This is because the bond operates at Layer 2 of the OSI model, the data link layer, and no IP addressing is used at this layer. IP addresses for virtual machines are associated with VIFs.

Bonding types

XenServer provides three different types of bonds, all of which can be configured using either the CLI or XenCenter:

- Active-Active mode, with VM traffic balanced between the bonded NICs. See [Active-active bonding](#).
- Active-Passive mode, where only one NIC actively carries traffic. See [Active-passive bonding](#).

- LACP Link Aggregation, in which active and stand-by NICs are negotiated between the switch and the host. See [LACP Link Aggregation Control Protocol bonding](#).

Note:

Bonding is set up with an Up Delay of 31,000 ms and a Down Delay of 200 ms. The seemingly long Up Delay is deliberate because of the time some switches take to enable the port. Without a delay, when a link comes back after failing, the bond can rebalance traffic onto it before the switch is ready to pass traffic. To move both connections to a different switch, move one, then wait 31 seconds for it to be used again before moving the other. For information about changing the delay, see [Changing the up delay for bonds](#).

Bond status

XenServer provides status for bonds in the event logs for each host. If one or more links in a bond fails or is restored, it is noted in the event log. Likewise, you can query the status of a bond's links by using the `links-up` parameter as shown in the following example:

```
1 xe bond-param-get uuid=bond_uuid param-name=links-up
```


XenServer checks the status of links in bonds approximately every five seconds. Therefore, if more links in the bond fail in the five-second window, the failure is not logged until the next status check.

Bonding event logs appear in the XenCenter **Notifications > Events** view. For users not running XenCenter, event logs also appear in `/var/log/xensource.log` on each host.

Active-active bonding

Active-active is an active/active configuration for guest traffic: both NICs can route VM traffic simultaneously. When bonds are used for management traffic, only one NIC in the bond can route traffic: the other NIC remains unused and provides failover support. Active-active mode is the default bonding mode.

You can bond either two, three, or four NICs in active-active mode. However, in active-active mode, bonding three, or four NICs is only beneficial for VM traffic, as shown in the illustration that follows.



Active-active bonds (vSwitch network stack)

XenServer can only send traffic over two or more NICs when there is more than one MAC address associated with the bond. XenServer can use the virtual MAC addresses in the VIF to send traffic across multiple links. Specifically:

- **VM traffic.** Provided you enable bonding on NICs carrying only VM (guest) traffic, all links are active and NIC bonding can balance spread VM traffic across NICs. An individual VIF's traffic is never split between NICs.

- **Management or storage traffic.** Only one of the links (NICs) in the bond is active and the other NICs remain unused unless traffic fails over to them. Configuring a management interface or secondary interface on a bonded network provides resilience.
- **Mixed traffic.** If the bonded NIC carries a mixture of IP-based storage traffic and guest traffic, only the guest and control domain traffic are load balanced. The control domain is essentially a virtual machine so it uses a NIC like the other guests. XenServer balances the control domain's traffic the same way as it balances VM traffic.

Traffic balancing XenServer balances the traffic between NICs by using the source MAC address of the packet. Because for management traffic, only one source MAC address is present, active-active mode can only use one NIC, and traffic is not balanced. Traffic balancing is based on two factors:

- The virtual machine and its associated VIF sending or receiving the traffic
- The quantity of data (in kilobytes) being sent.

XenServer evaluates the quantity of data (in kilobytes) each NIC is sending and receiving. If the quantity of data sent across one NIC exceeds the quantity of data sent across the other NIC, XenServer rebalances which VIFs use which NICs. The VIF's entire load is transferred. One VIF's load is never split between two NICs.

Though active-active NIC bonding can provide load balancing for traffic from multiple VMs, it cannot provide a single VM with the throughput of two NICs. Any given VIF only uses one of the links in a bond at a time. As XenServer periodically rebalances traffic, VIFs are not permanently assigned to a specific NIC in the bond.

Active-active mode is sometimes described as Source Load Balancing (SLB) bonding as XenServer uses SLB to share load across bonded network interfaces. SLB is derived from the open-source Adaptive Load Balancing (ALB) mode and reuses the ALB functionality to rebalance load across NICs dynamically.

When rebalancing, the number of bytes going over each secondary (interface) is tracked over a given period. If a packet to be sent contains a new source MAC address, it is assigned to the secondary interface with the lowest utilization. Traffic is rebalanced at regular intervals.

Each MAC address has a corresponding load and XenServer can shift entire loads between NICs depending on the quantity of data a VM sends and receives. For active-active traffic, all the traffic from one VM can be sent on only one NIC.

Note:


Active-active bonding does not require switch support for EtherChannel or 802.3ad (LACP).

Active-passive bonding

An active-passive bond routes traffic over only one of the NICs. If the active NIC loses network connectivity, traffic fails over to the other NIC in the bond. Active-passive bonds route traffic over the active NIC. The traffic shifts to the passive NIC if the active NIC fails.

Active-passive bonding can bond two, three, or four NICs together. However, regardless of the traffic type, when you bond NICs in active-passive mode, only one link is active and there is no load balancing between links.

The illustration that follows shows two bonded NICs configured in active-passive mode.



Active-active mode is the default bonding configuration in XenServer. If you are configuring bonds using the CLI, you must specify a parameter for the active-passive mode. Otherwise, an active-active bond is created. You do not need to configure active-passive mode because a network is carrying management traffic or storage traffic.

Active-passive can be a good choice for resiliency as it offers several benefits. With active-passive bonds, traffic does not move around between NICs. Similarly, active-passive bonding lets you configure two switches for redundancy but does not require stacking. If the management switch dies, stacked switches can be a single point of failure.

Active-passive mode does not require switch support for EtherChannel or 802.3ad (LACP).

Consider configuring active-passive mode in situations when you do not need load balancing or when you only intend to send traffic on one NIC.

Important:

After you have created VIFs or your pool is in production, be careful about changing bonds or creating bonds.

LACP Link Aggregation Control Protocol bonding

LACP Link Aggregation Control Protocol is a type of bonding that bundles a group of ports together and treats it like a single logical channel. LACP bonding provides failover and can increase the total amount of bandwidth available.

Unlike other bonding modes, LACP bonding requires configuring both sides of the links: creating a bond on the host, and creating a Link Aggregation Group (LAG) for each bond on the switch. See [Switch configuration for LACP bonds](#). Your switches must support the IEEE 802.3ad standard to use LACP bonding.

A comparison of active-active SLB bonding and LACP bonding:

Active-active SLB bonding Benefits:

- Can be used with any switch on the Hardware Compatibility List.
- Does not require switches that support stacking.
- Supports four NICs.

Considerations:

- Optimal load balancing requires at least one NIC per VIF.
- Storage or management traffic cannot be split on multiple NICs.
- Load balancing occurs only if multiple MAC addresses are present.

LACP bonding Benefits:

- All links can be active regardless of traffic type.
- Traffic balancing does not depend on source MAC addresses, so all traffic types can be balanced.

Considerations:

- Switches must support the IEEE 802.3ad standard.
- Requires switch-side configuration.
- Requires a single switch or stacked switch.

Traffic balancing XenServer supports two LACP bonding hashing types. The term hashing describes how the NICs and the switch distribute the traffic—(1) load balancing based on IP and port of source and destination addresses and (2) load balancing based on source MAC address.

Depending on the hashing type and traffic pattern, LACP bonding can potentially distribute traffic more evenly than active-active NIC bonding.

Note:


You configure settings for outgoing and incoming traffic separately on the host and the switch: the configuration does not have to match on both sides.

Load balancing based on IP and port of source and destination addresses.


This hashing type is the default LACP bonding hashing algorithm. If there is a variation in the source or destination IP or port numbers, traffic from one guest can be distributed over two links.

If a virtual machine is running several applications which use different IP or port numbers, this hashing type distributes traffic over several links. Distributing the traffic gives the guest the possibility of using the aggregate throughput. This hashing type lets one guest use the whole throughput of multiple NICs.

As shown in the illustration that follows, this hashing type can distribute the traffic of two different applications on a virtual machine to two different NICs.



Configuring LACP bonding based on IP and port of source and destination address is beneficial when you want to balance the traffic of two different applications on the same VM. For example, when only one virtual machine is configured to use a bond of three NICs.




Load balancing based on IP and port of source and destination addresses

The balancing algorithm for this hashing type uses five factors to spread traffic across the NICs: the source IP address, source port number, destination IP address, destination port number, and source MAC address.

Load balancing based on source MAC address.

This type of load balancing works well when there are multiple virtual machines on the same host. Traffic is balanced based on the virtual MAC address of the VM from which the traffic originated. XenServer sends outgoing traffic using the same algorithm as it does in active-active bonding. Traffic coming from the same guest is not split over multiple NICs. As a result, this hashing type is not suitable if there are fewer VIFs than NICs: load balancing is not optimal because the traffic cannot be split across NICs.



Switch configuration

Depending on your redundancy requirements, you can connect the NICs in the bond to either the same or separate stacked switches. If you connect one of the NICs to a second, redundant switch and a NIC or switch fails, traffic fails over to the other NIC. Adding a second switch prevents a single point-of-failure in your configuration in the following ways:

- When you connect one of the links in a bonded management interface to a second switch, if the switch fails, the management network remains online and the hosts can still communicate with each other.
- If you connect a link (for any traffic type) to a second switch and the NIC or switch fails, the virtual machines remain on the network as their traffic fails over to the other NIC/switch.

Use stacked switches when you want to connect bonded NICs to multiple switches and you configured the LACP bonding mode. The term ‘stacked switches’ is used to describe configuring multiple

physical switches to function as a single logical switch. You must join the switches together physically and through the switch-management software so the switches function as a single logical switching unit, as per the switch manufacturer's guidelines. Typically, switch stacking is only available through proprietary extensions and switch vendors may market this functionality under different terms.

Note:

If you experience issues with active-active bonds, the use of stacked switches may be necessary. Active-passive bonds do not require stacked switches.

Switch configuration for LACP bonds Because the specific details of switch configuration vary by manufacturer, there are a few key points to remember when configuring switches for use with LACP bonds:

- The switch must support LACP and the IEEE 802.3ad standard.
- When you create the LAG group on the switch, you must create one LAG group for each LACP bond on the host. For example, if you have a five-host pool and you created an LACP bond on NICs 4 and 5 on each host, you must create five LAG groups on the switch. One group for each set of ports corresponding with the NICs on the host.

You may also need to add your VLAN ID to your LAG group.

- XenServer LACP bonds require setting the Static Mode setting in the LAG group to be set to Disabled.

As previously mentioned in *Switch configuration*, stacking switches are required to connect LACP bonds to multiple switches.

Initial networking configuration after setup

The XenServer host networking configuration is specified during initial host installation. Options such as IP address configuration (DHCP/static), the NIC used as the management interface, and host name are set based on the values provided during installation.

When a host has multiple NICs, the configuration present after installation depends on which NIC is selected for management operations during installation:

- PIFs are created for each NIC in the host
- The PIF of the NIC selected for use as the management interface is configured with the IP addressing options specified during installation
- A network is created for each PIF ("network 0", "network 1", and so on)
- Each network is connected to one PIF

- The IP addressing options are left unconfigured for all PIFs other than the PIF used as the management interface

When a host has a single NIC, the following configuration is present after installation:

- A single PIF is created corresponding to the host's single NIC
- The PIF is configured with the IP addressing options specified during installation and to enable management of the host
- The PIF is set for use in host management operations
- A single network, network 0, is created
- Network 0 is connected to the PIF to enable external connectivity to VMs

When an installation of XenServer is done on a tagged VLAN network, the following configuration is present after installation:

- PIFs are created for each NIC in the host
- The PIF for the tagged VLAN on the NIC selected for use as the management interface is configured with the IP address configuration specified during installation
- A network is created for each PIF (for example: network 1, network 2, and so on).
An additional VLAN network is created (for example, for Pool-wide network associated with the first network interface on VLAN<TAG>)
- Each network is connected to one PIF. The VLAN PIF is set for use in host management operations

In both cases, the resulting networking configuration allows connection to the XenServer host by XenCenter, the xe CLI, and any other management software running on separate machines through the IP address of the management interface. The configuration also provides external networking for VMs created on the host.

The PIF used for management operations is the only PIF ever configured with an IP address during XenServer installation. External networking for VMs is achieved by bridging PIFs to VIFs using the network object which acts as a virtual Ethernet switch.

The steps required for networking features such as VLANs, NIC bonds, and dedicating a NIC to storage traffic are covered in the sections that follow.

Changing networking configuration

You can change your networking configuration by modifying the network object. To do so, you run a command that affects either the network object or the VIF.

Modifying the network object

You can change aspects of a network, such as the frame size (MTU), name-label, name-description, purpose, and other values. Use the `xe network-param-set` command and its associated parameters to change the values.

When you run the `xe network-param-set` command, the only required parameter is `uuid`.

Optional parameters include:

- `default_locking_mode`. See [Simplifying VIF locking mode configuration in the Cloud](#).
- `name-label`
- `name-description`
- `MTU`
- `purpose`. See [Adding a purpose to a network](#).
- `other-config`

If a value for a parameter is not given, the parameter is set to a null value. To set a (key, value) pair in a map parameter, use the syntax `map-param:key=value`.

Changing the up delay for bonds

Bonding is set up with an Up Delay of 31,000 ms by default to prevent traffic from being rebalanced onto a NIC after it fails. While seemingly long, the up delay is important for all bonding modes and not just active-active.

However, if you understand the appropriate settings to select for your environment, you can change the up delay for bonds by using the procedure that follows.

Set the up delay in milliseconds:

```
1 xe pif-param-set uuid=<uuid of bond interface PIF> other-config:bond-updelay=<delay in ms>
```

To make the change take effect, you must unplug and then replug the physical interface:

```
1 xe pif-unplug uuid=<uuid of bond interface PIF>
```

```
1 xe pif-plug uuid=<uuid of bond interface PIF>
```

Manage networking

June 22, 2026

Network configuration procedures in this section differ depending on whether you are configuring a stand-alone host or a host that is part of a resource pool.

Create networks in a standalone host

Because external networks are created for each PIF during host installation, creating extra networks is typically only required to:

- Use a private network
- Support advanced operations such as VLANs or NIC bonding

For information about how to add or delete networks using XenCenter, see [Add a New Network](#) in the XenCenter documentation.

Open the XenServer® host text console.

Create the network by using the `network-create` command, which returns the UUID of the newly created network:

```
1 xe network-create name-label=mynetwork
```

At this point, the network is not connected to a PIF and therefore is internal.

Create networks in resource pools

All XenServer hosts in a resource pool must have the same number of physical NICs. This requirement is not strictly enforced when a host is joined to a pool. One of the NICs is always designated as the *management interface*, used for XenServer management traffic.

As all hosts in a pool share a common set of networks, it is important to have the same physical networking configuration for XenServer hosts in a pool.

PIFs on the individual hosts are connected to pool-wide networks based on device name.

Network X corresponds to the network interface in position X under the [Dom0 network naming model](#).

If one XenServer host has a different number of NICs than other hosts in the pool, complications can arise. The complications can arise because not all pool networks are valid for all pool hosts. For example, if hosts *host1* and *host2* are in the same pool and *host1* has four NICs and *host2* only has two, only the networks connected to PIFs corresponding to the first two network interfaces from *host1* are

valid on *host2*. VMs on *host1* with VIFs connected to networks corresponding to the remaining network interfaces cannot migrate to host *host2*.

Create VLANs

For hosts in a resource pool, you can use the `pool-vlan-create` command. This command creates the VLAN and automatically creates and plug-ins the required PIFs on the hosts in the pool. For more information, see `pool-vlan-create`.

Open the XenServer host console.

Create a network for use with the VLAN. The UUID of the new network is returned:

```
1 xe network-create name-label=network5
```

Use the `pif-list` command to find the UUID of the PIF corresponding to the physical NIC supporting the desired VLAN tag. The UUIDs and device names of all PIFs are returned, including any existing VLANs:

```
1 xe pif-list
```

Create a VLAN object specifying the desired physical PIF and VLAN tag on all VMs to be connected to the new VLAN. A new PIF is created and plugged to the specified network. The UUID of the new PIF object is returned.

```
1 xe vlan-create network-uuid=network_uuid pif-uuid=pif_uuid vlan=5
```

Attach VM VIFs to the new network. For more information, see [Creating networks in a standalone host](#).

Create NIC bonds on a standalone host

We recommend using XenCenter to create NIC bonds. For more information, see [Configuring NICs](#).

This section describes how to use the `xe` CLI to bond NIC interfaces on XenServer hosts that are not in a pool. For information on using the `xe` CLI to create NIC bonds on XenServer hosts that comprise a resource pool, see *Creating NIC bonds in resource pools*.

Create a NIC bond

When you bond a NIC, the bond absorbs the PIF/NIC in use as the management interface. The management interface is automatically moved to the bond PIF.

1. Use the `network-create` command to create a network for use with the bonded NIC. The UUID of the new network is returned:

```
1 xe network-create name=label=bond0
```

2. Use the `pif-list` command to determine the UUIDs of the PIFs to use in the bond:

```
1 xe pif-list
```

3. Do one of the following:

- To configure the bond in active-active mode (default), use the `bond-create` command to create the bond. Using commas to separate the parameters, specify the newly created network UUID and the UUIDs of the PIFs to be bonded:

```
1 xe bond-create network-uuid=network_uuid /  
2     pif-uuids=pif_uuid_1,pif_uuid_2,pif_uuid_3,pif_uuid_4
```

Type two UUIDs when you are bonding two NICs and four UUIDs when you are bonding four NICs. The UUID for the bond is returned after running the command.

- To configure the bond in active-passive or LACP bond mode, use the same syntax, add the optional `mode` parameter, and specify `lacp` or `active-backup`:

```
1 xe bond-create network-uuid=network_uuid pif-uuids=pif_uuid_1  
    , /  
2     pif_uuid_2,pif_uuid_3,pif_uuid_4 /  
3     mode=balance-slb | active-backup | lacp
```

Control the MAC address of the bond

When you bond the management interface, it subsumes the PIF/NIC in use as the management interface. If the host uses DHCP, the bond's MAC address is the same as the PIF/NIC in use. The management interface's IP address can remain unchanged.

You can change the bond's MAC address so that it is different from the MAC address for the (current) management-interface NIC. However, as the bond is enabled and the MAC/IP address in use changes, existing network sessions to the host are dropped.

You can control the MAC address for a bond in two ways:

- An optional `mac` parameter can be specified in the `bond-create` command. You can use this parameter to set the bond MAC address to any arbitrary address.
- If the `mac` parameter is not specified, XenServer uses the MAC address of the management interface if it is one of the interfaces in the bond. If the management interface is not part of the bond, but another management interface is, the bond uses the MAC address (and also the IP address)

of that management interface. If none of the NICs in the bond is a management interface, the bond uses the MAC of the first named NIC.

Revert NIC bonds

When reverting the XenServer host to a non-bonded configuration, the `bond-destroy` command automatically configures the primary NIC as the interface for the management interface. Therefore, all VIFs are moved to the management interface. If management interface of a host is on tagged VLAN bonded interface, on performing `bond-destroy`, management VLAN is moved to primary NIC.

The term primary NIC refers to the PIF that the MAC and IP configuration was copied from when creating the bond. When bonding two NICs, the primary NIC is:

1. The management interface NIC (if the management interface is one of the bonded NICs).
2. Any other NIC with an IP address (if the management interface was not part of the bond).
3. The first named NIC. You can find out which one it is by running the following:

```
1 xe bond-list params=all
```

Create NIC bonds in resource pools

Whenever possible, create NIC bonds as part of initial resource pool creation, before joining more hosts to the pool or creating VMs. Doing so allows the bond configuration to be automatically replicated to hosts as they are joined to the pool and reduces the number of steps required.

Adding a NIC bond to an existing pool requires one of the following:

- Using the CLI to configure the bonds on the pool coordinator and then each member of the pool.
- Using the CLI to configure bonds on the pool coordinator and then restarting each pool member so that it inherits its settings from the pool coordinator.
- Using XenCenter® to configure the bonds on the pool coordinator. XenCenter automatically synchronizes the networking settings on the member hosts with the pool coordinator, so you do not need to restart the member hosts.

For simplicity and to prevent misconfiguration, we recommend using XenCenter to create NIC bonds. For more information, see [Configuring NICs](#).

This section describes using the `xe` CLI to create bonded NIC interfaces on XenServer hosts that comprise a resource pool. For information on using the `xe` CLI to create NIC bonds on a standalone host, see *Creating NIC bonds on a standalone host*.

Warning:

Do not attempt to create network bonds when high availability is enabled. The process of bond creation disturbs the in-progress high availability heartbeat and causes hosts to self-fence (shut themselves down). The hosts can fail to restart properly and may need the `host-emergency-ha-disable` command to recover.

Select the host you want to be the pool coordinator. The pool coordinator belongs to an unnamed pool by default. To create a resource pool with the CLI, rename the existing nameless pool:

```
1 xe pool-param-set name-label="New Pool" uuid=pool_uuid
```

Create the NIC bond as described in [Create a NIC bond](#).

Open a console on a host that you want to join to the pool and run the command:

```
1 xe pool-join master-address=host1 master-username=root master-password=
  password
```

The network and bond information is automatically replicated to the new host. The management interface is automatically moved from the host NIC where it was originally configured to the bonded PIF. That is, the management interface is now absorbed into the bond so that the entire bond functions as the management interface.

Use the `host-list` command to find the UUID of the host being configured:

```
1 xe host-list
```

Warning:

Do not attempt to create network bonds while high availability is enabled. The process of bond creation disturbs the in-progress high availability heartbeat and causes hosts to self-fence (shut themselves down). The hosts can fail to restart properly and you may need to run the `host-emergency-ha-disable` command to recover.

Configure a dedicated storage NIC

You can use XenCenter or the xe CLI to assign a NIC an IP address and dedicate it to a specific function, such as storage traffic. When you configure a NIC with an IP address, you do so by creating a secondary interface. (The IP-enabled NIC XenServer used for management is known as the management interface.)

When you want to dedicate a secondary interface for a specific purpose, ensure that the appropriate network configuration is in place. This is to ensure that the NIC is used only for the desired traffic. To dedicate a NIC to storage traffic, configure the NIC, storage target, switch, and VLAN such that the

target is only accessible over the assigned NIC. If your physical and IP configuration does not limit the traffic sent across the storage NIC, you can send traffic, such as management traffic across the secondary interface.

When you create a new secondary interface for storage traffic, you must assign it an IP address that is:

- On the same subnet as the storage controller, if applicable, and
- Not on the same subnet as any other secondary interfaces or the management interface.

When you are configuring secondary interfaces, each secondary interface must be on a separate subnet. For example, if you want to configure two more secondary interfaces for storage, you require IP addresses on three different subnets –one subnet for the management interface, one subnet for Secondary Interface 1, and one subnet for Secondary Interface 2.

Note:

When selecting a NIC to configure as a secondary interface for use with iSCSI or NFS SRs, ensure that the dedicated NIC uses a separate IP subnet that is not routable from the management interface. If this is not enforced, then storage traffic may be directed over the main management interface after a host restart, because of the order in which network interfaces are initialized.

Ensure that the PIF is on a separate subnet, or routing is configured to suit your network topology to force desired traffic over the selected PIF.

Set up an IP configuration for the PIF, adding appropriate values for the mode parameter. If using static IP addressing, add the IP, netmask, gateway, and DNS parameters:

```
1 xe pif-reconfigure-ip mode=DHCP | Static uuid=pif-uuid
```

Set the PIF's disallow-unplug parameter to true:

```
1 xe pif-param-set disallow-unplug=true uuid=pif-uuid
```

```
1 xe pif-param-set other-config:management_purpose="Storage" uuid=pif-uuid
```

If you want to use a secondary interface for storage that can be routed from the management interface also (bearing in mind that this configuration is not the best practice), you have two options:

- After a host restart, ensure that the secondary interface is correctly configured. Use the `xe pbd-unplug` and `xe pbd-plug` commands to reinitialize the storage connections on the host. This command restarts the storage connection and routes it over the correct interface.
- Alternatively, you can use `xe pif-forget` to delete the interface from the XenServer database and manually configure it in the control domain. `xe pif-forget` is an advanced option and requires you to be familiar with how to configure Linux networking manually.

Use SR-IOV enabled NICs

Single Root I/O Virtualization (SR-IOV) is a virtualization technology that allows a single PCI device to appear as multiple PCI devices on the physical system. The actual physical device is known as a Physical Function (PF) while the others are known as Virtual Functions (VF). The hypervisor can assign one or more VFs to a Virtual Machine (VM): the guest can then use the device as if it were directly assigned.

Assigning one or more NIC VFs to a VM allows its network traffic to bypass the virtual switch. When configured, each VM behaves as though it is using the NIC directly, reducing processing overhead, and improving performance.

Benefits of SR-IOV

An SR-IOV VF has a better performance than VIF. It can ensure the hardware-based segregation between traffic from different VMs through the same NIC (bypassing the XenServer network stack).

Using this feature, you can:

- Enable SR-IOV on NICs that support SR-IOV.
- Disable SR-IOV on NICs that support SR-IOV.
- Manage SR-IOV VFs as a VF resource pool.
- Assign SR-IOV VFs to a VM.
- Configure SR-IOV VFs (For example, MAC address, VLAN, rate).
- Run tests to confirm if SR-IOV is supported as part of the Automated Certification Kit.

System configuration

Configure the hardware platform correctly to support SR-IOV. The following technologies are required:

- I/O MMU virtualization (AMD-Vi and Intel VT-d)
- Alternative Routing-ID Interpretation (ARI)
- Address Translation Services (ATS)
- Access Control Services (ACS)

Check the documentation that comes with your system for information on how to configure the system firmware to enable the mentioned technologies.

Enable an SR-IOV network on a NIC

In XenCenter, use the **New Network** wizard in the **Networking** tab to create and enable an SR-IOV network on a NIC.

Assign an SR-IOV network to the virtual interface (VM level)

In XenCenter, at the VM level, use the **Add Virtual Interface** wizard in the **Networking** tab to add an SR-IOV enabled network as a virtual interface for that VM. For more information, see [Add a New Network](#).

Supported NICs and guests

For a list of supported hardware platforms and NICs, see [Hardware Compatibility List](#). See the documentation provided by the vendor for a particular guest to determine whether it supports SR-IOV.

Limitations

- For certain NICs using legacy drivers (for example, Intel I350 family) the host must be rebooted to enable or disable SR-IOV on these devices.
- A pool level SR-IOV network having different types of NICs are not supported.
- An SR-IOV VF and a normal VIF from the same NIC may not be able to communicate with each other because of the NIC hardware limitations. To enable these VMs to communicate, ensure that communication uses the pattern VF to VF or VIF to VIF, and not VF to VIF.
- Quality of Service settings for some SR-IOV VFs do not take effect because they do not support network speed rate limiting.
- Performing live migration, suspend, and checkpoint is not supported on VMs using an SR-IOV VF.
- SR-IOV VFs do not support hot-plugging.
- SR-IOV VFs do not support network boot.
- For some NICs with legacy NIC drivers, rebooting may be required even after host restart which indicates that the NIC is not able to enable SR-IOV.
- If your VM has an SR-IOV VF, functions that require Live Migration are not possible. This is because the VM is directly tied to the physical SR-IOV enabled NIC VF.

- SR-IOV can be used in an environment that makes use of high availability. However, SR-IOV is not considered in the capacity planning. VMs that have SR-IOV VFs assigned are restarted on a best-effort basis when there is a host in the pool that has appropriate resources. These resources include SR-IOV enabled on the right network and a free VF.
- SR-IOV VFs are not supported with the PVS-Accelerator.

Configure SR-IOV VFs for legacy drivers

Usually the maximum number of VFs that a NIC can support can be determined automatically. For NICs using legacy drivers (for example, Intel I350 family), the limit is defined within the driver module configuration file. The limit may need to be adjusted manually. To set it to the maximum, open the file using an editor and change the line starting:

```
1 ## VFs-maxvfs-by-user:
```

For example, to set the maximum VFs to 4 for the `igb` driver edit `/etc/modprobe.d/igb.conf` to read:

```
1 ## VFs-param: max_vfs
2 ## VFs-maxvfs-by-default: 7
3 ## VFs-maxvfs-by-user: 4
4 options igb max_vfs=0
```

Notes:

- The value must be less than or equal to the value in the line `VFs-maxvfs-by-default`.
- Do not change any other line in these files.
- Make the changes before enabling SR-IOV.

CLI

See [SR-IOV commands](#) for CLI instructions on creating, deleting, displaying SR-IOV networks and assigning an SR-IOV VF to a VM.

Control the rate of outgoing data (QoS)

To limit the amount of *outgoing* data a VM can send per second, set an optional Quality of Service (QoS) value on VM virtual interfaces (VIFs). The setting lets you specify a maximum transmit rate for outgoing packets in *kilobytes* per second.

The Quality of Service value limits the rate of transmission *from* the VM. The Quality of Service setting does not limit the amount of data the VM can receive. If such a limit is desired, we recommend limiting the rate of incoming packets higher up in the network (for example, at the switch level).

Depending on networking stack configured in the pool, you can set the Quality of Service value on VM virtual interfaces (VIFs) in one of two places. You can set this value either by using the xe CLI or in XenCenter.

- **XenCenter** You can set the Quality of Service transmit rate limit value in the properties dialog for the virtual interface.
- **xe commands** You can set the Quality of Service transmit rate using the CLI using the commands in the section that follow.

Example of CLI command for QoS

To limit a VIF to a maximum transmit rate of 100 kilobytes per second using the CLI, use the `vif-param-set` command:

```
1 xe vif-param-set uuid=vif_uuid qos_algorithm_type=ratelimit
2 xe vif-param-set uuid=vif_uuid qos_algorithm_params:kbps=100
```

Note:

The `kbps` parameter denotes *kilobytes* per second (kBps), not kilobits per second (kbps).

Change networking configuration options

This section discusses how to change the networking configuration of your XenServer host. It includes:

- Changing the hostname (that is, the Domain Name System (DNS) name)
- Adding or deleting DNS servers
- Changing IP addresses
- Changing which NIC is used as the management interface
- Adding a new physical NIC to the server
- Adding a purpose to a network
- Enabling ARP filtering (switch-port locking)

Hostname

The system hostname, also known as the domain or DNS name, is defined in the pool-wide database and changed using the `xe host-set-hostname-live` CLI command as follows:

```
1 xe host-set-hostname-live host-uuid=host_uuid host-name=host-name
```

The underlying control domain hostname changes dynamically to reflect the new hostname.

DNS servers

To add or delete DNS servers in the IP addressing configuration of the XenServer host, use the `pif-reconfigure-ip` command. For example, for a PIF with a static IP:

```
1 xe pif-reconfigure-ip uuid=pif_uuid mode=static DNS=new_dns_ip IP=IP
  netmask=netmask
```

Change IP address configuration for a standalone host

You can use the `xe` CLI to change the network interface configuration. Do not change the underlying network configuration scripts directly.

To change the IP address configuration of a PIF, use the `pif-reconfigure-ip` CLI command. See `pif-reconfigure-ip` for details on the parameters of the `pif-reconfigure-ip` command. See the following section for information on changing host IP addresses in resource pools.

Change IP address configuration in resource pools

XenServer hosts in resource pools have a single management IP address used for management and communication to and from other hosts in the pool. The steps required to change the IP address of a host's management interface are different for pool coordinator and other hosts.

Note:

You must be careful when changing the IP address of a host, and other networking parameters. Depending upon the network topology and the change being made, connections to network storage can be lost. When this happens, the storage must be replugged using the **Repair Storage** function in XenCenter, or by using the `pbid-plug` CLI command. For this reason, we recommend that you migrate VMs away from the host before changing its IP configuration.

Use the `pif-reconfigure-ip` CLI command to set the IP address as desired. See `pif-reconfigure-ip` for details on the parameters of the `pif-reconfigure-ip` command.
:

```
1 xe pif-reconfigure-ip uuid=pif_uuid mode=DHCP
```

Use the `host-list` CLI command to confirm that the member host has successfully reconnected to the pool coordinator by checking that all the other XenServer hosts in the pool are visible:

```
1 xe host-list
```

Changing the IP address of the pool coordinator XenServer host requires extra steps. This is because each pool member uses the advertised IP address of the pool coordinator for communication. The pool members do not know how to contact the pool coordinator when its IP address changes.

Whenever possible, use a dedicated IP address that is not likely to change for the lifetime of the pool for pool coordinators.

Use the `pif-reconfigure-ip` CLI command to set the IP address as desired:

```
1 xe pif-reconfigure-ip uuid=pif_uuid mode=DHCP
```

When the IP address of the pool coordinator changes, all member hosts enter into an emergency mode when they fail to contact the pool coordinator.

On the pool coordinator, use the `pool-recover-slaves` command to force the pool coordinator to contact each pool member and inform them of the new pool coordinator IP address:

```
1 xe pool-recover-slaves
```

Management interface

When you install XenServer on a host, one of its NICs is designated as the *management interface*: the NIC used for XenServer management traffic. The management interface is used for XenCenter connections to the host (for example, Citrix Virtual Apps and Desktops™) and for host-to-host communication.

Use the `pif-list` command to determine which PIF corresponds to the NIC to be used as the management interface. The UUID of each PIF is returned.

```
1 xe pif-list
```

Use the `pif-param-list` command to verify the IP addressing configuration for the PIF used for the management interface. If necessary, use the `pif-reconfigure-ip` command to configure IP addressing for the PIF to be used.

```
1 xe pif-param-list uuid=pif_uuid
```

Use the `host-management-reconfigure` CLI command to change the PIF used for the management interface. If this host is part of a resource pool, *this command must be issued on the member host*

console:

```
1 xe host-management-reconfigure pif-uuid=pif_uuid
```

Use the `network-list` command to determine which PIF corresponds to the NIC to be used as the management interface for all the hosts in the pool. The UUID of pool wide network is returned.

```
1 xe network-list
```

Use the `network-param-list` command to fetch the PIF UUIDs of all the hosts in the pool. Use the `pif-param-list` command to verify the IP addressing configuration for the PIF for the management interface. If necessary, use the `pif-reconfigure-ip` command to configure IP addressing for the PIF to be used.

```
1 xe pif-param-list uuid=pif_uuid
```

Use the `pool-management-reconfigure` CLI command to change the PIF used for the management interface listed in the Networks list.

```
1 xe pool-management-reconfigure network-uuid=network_uuid
```

Restrict use of port 80

You can use either HTTPS over port 443 or HTTP over port 80 to communicate with XenServer. For security reasons, by default port 80 is closed for new installations. Upgrades retain existing port settings.

To open or close port 80, see the `https-only` `xe` CLI command or [Change Pool Properties](#) in the XenCenter documentation.

Disable management access

To disable remote access to the management console entirely, use the `host-management-disable` CLI command.

Warning:

When the management interface is disabled, you must log in on the physical host console to perform management tasks. External interfaces such as XenCenter do not work when the management interface is disabled.

Add a new physical NIC

1. Install a new physical NIC on your XenServer host in the usual manner.

2. Restart your XenServer host.
3. List all the physical NICs for that XenServer host by using the following command:

```
1 xe pif-list host-uuid=<host_uuid>
```

4. If you do not see the additional NIC, scan for new physical interfaces by using the following command:

```
1 xe pif-scan host-uuid=<host_uuid>
```

This command creates a new PIF object for the new NIC.

5. List the physical NICs on the XenServer host again to verify that the new NIC is visible:

```
1 xe pif-list host-uuid=<host_uuid>
```

6. The new PIF is initially listed as disconnected (`currently-attached ( R0): false`). To bring it up, use the following command:

```
1 xe pif-plug uuid=<uuid_of_pif>
```

Alternatively, you can use XenCenter to rescan for new NICs. For more information, see [Configuring NICs](#) in the XenCenter documentation.

Remove a physical NIC

Before removing the NIC, ensure that you know the UUID of the corresponding PIF. Remove the physical NIC from your XenServer host in the usual manner. After restarting the host, run the `xe` CLI command `pif-forget uuid=<UUID>` to destroy the PIF object.

Add a purpose to a network

The network purpose can be used to add extra functionalities to a network. For example, the ability to use the network to make NBD connections.

To add a network purpose, use the `xe network-param-add` command:

```
1 xe network-param-add param-name=purpose param-key=purpose uuid=network-  
  uuid
```

To delete a network purpose, use the `xe network-param-remove` command:

```
1 xe network-param-remove param-name=purpose param-key=purpose uuid=  
  network-uuid
```

Currently, the available values for the network purpose are `nbd` and `insecure_nbd`. For more information, see the [XenServer Changed Block Tracking Guide](#).

Use switch port locking

The XenServer switch-port locking feature lets you control traffic sent from unknown, untrusted, or potentially hostile VMs by limiting their ability to pretend they have a MAC or IP address that was not assigned to them. You can use the port-locking commands to block all traffic on a network by default or define specific IP addresses from which an individual VM is allowed to send traffic.

Using switch-port locking lets you simplify your network configuration by enabling all of your tenants or guests to use the same Layer 2 network.

One of the most important functions of the port-locking commands is they can restrict the traffic that an untrusted guest send. This restricts the guest's ability to pretend it has a MAC or IP address it does not actually possess. Specifically, you can use these commands to prevent a guest from:

- Claiming an IP or MAC address other than the ones the XenServer administrator has specified it can use
- Intercepting, spoofing, or disrupting the traffic of other VMs

Requirements

- When you enable Role Based Access Control (RBAC) in your environment, the user configuring switch-port locking must be logged in with an account that has at least a Pool Operator or Pool Admin role. When RBAC is not enabled in your environment, the user must be logged in with the root account for the pool coordinator.
- When you run the switch-port locking commands, networks can be online or offline.
- In Windows guests, the disconnected Network icon only appears when XenServer VM Tools are installed in the guest.

Notes Without any switch-port locking configurations, VIFs are set to “network_default” and Networks are set to “unlocked.”

Configuring switch-port locking is not supported when any third-party controllers are in use in the environment.

Switch port locking does not prevent cloud tenants from:

- Performing an IP-level attack on another tenant/user. However, switch-port locking prevents them performing the IP-level attack if they attempt to use the following means to do so and switch-port locking is configured: a) impersonating another tenant in the cloud or user or b) initiating an intercept of traffic intended for another user.
- Exhausting network resources.

- Receiving some traffic intended for other virtual machines through normal switch flooding behaviors (for broadcast MAC addresses or unknown destination MAC addresses).

Likewise, switch-port locking does not restrict where a VM can send traffic to.

Implementation notes You can implement the switch-port locking functionality either by using the command line or the XenServer API. However, in large environments, where automation is a primary concern, the most typical implementation method might be by using the API.

Examples This section provides examples of how switch-port locking can prevent certain types of attacks. In these examples, VM-c is a virtual machine that a hostile tenant (Tenant C) is leasing and using for attacks. VM-a and VM-b are virtual machines leased by non-attacking tenants.

Example 1: How switch port locking can prevent ARP spoofing prevention:

ARP spoofing is used to indicate an attacker's attempts to associate their MAC address with the IP address for another node. ARP spoofing can potentially result in the node's traffic being sent to the attacker instead. To achieve this goal the attacker sends fake (spoofed) ARP messages to an Ethernet LAN.

Scenario:

Virtual Machine A (VM-a) wants to send IP traffic from VM-a to Virtual Machine B (VM-b) by addressing it to VM-b's IP address. The owner of Virtual Machine C wants to use ARP spoofing to pretend their VM, VM-c, is actually VM-b.

1. VM-c sends a speculative stream of ARP replies to VM-a. The ARP replies claim that the MAC address in the reply (c_MAC) is associated with the IP address, b_IP

Result: Because the administrator enabled switch-port locking, these packets are all dropped because enabling switch-port locking prevents impersonation.

2. VM-b sends an ARP reply to VM-a, claiming that the MAC address in the reply (b_MAC) is associated with the IP address, b_IP.

Result: VM-a receives VM-b's ARP response.

Example 2: IP Spoofing prevention:

IP address spoofing is a process that conceals the identity of packets by creating Internet Protocol (IP) packets with a forged source IP address.

Scenario:

Tenant C is attempting to perform a Denial of Service attack using their host, Host-C, on a remote system to disguise their identity.

Attempt 1:

Tenant C sets Host-C's IP address and MAC address to VM-a's IP and MAC addresses (a_IP and a_MAC). Tenant C instructs Host-C to send IP traffic to a remote system.

Result: The Host-C packets are dropped. This is because the administrator enabled switch-port locking. The Host-C packets are dropped because enabling switch-port locking prevents impersonation.

Attempt 2:

Tenant C sets Host-C's IP address to VM-a's IP address (a_IP) and keeps their original c_MAC.

Tenant C instructs Host-C to send IP traffic to a remote system.

Result: The Host-C packets are dropped. This is because the administrator enabled switch-port locking, which prevents impersonation.

Example 3: Web hosting:

Scenario:

Alice is an infrastructure administrator.


One of her tenants, Tenant B, is hosting multiple websites from their VM, VM-b. Each website needs a distinct IP address hosted on the same virtual network interface (VIF).

Alice reconfigures Host-B's VIF to be locked to a single MAC but many IP addresses.

How switch-port locking works The switch-port locking feature lets you control packet filtering at one or more of two levels:

- **VIF level.** Settings you configure on the VIF determine how packets are filtered. You can set the VIF to prevent the VM from sending any traffic, restrict the VIF so it can only send traffic using its assigned IP address, or allow the VM to send traffic to any IP address on the network connected to the VIF.
- **Network level.** The XenServer network determines how packets are filtered. When a VIF's locking mode is set to `network_default`, it refers to the network-level locking setting to determine what traffic to allow.

VIF locking-mode states The XenServer switch-port locking feature provides a locking mode that lets you configure VIFs in four different states. These states only apply when the VIF is plugged into a running virtual machine.



- **Network_default.** When the VIF's state is set to `network_default`, XenServer uses the network's `default-locking-mode` parameter to determine if and how to filter packets traveling through the VIF. The behavior varies according to if the associated network has the network default locking mode parameter set to disabled or unlocked:

-`default-locking-mode=disabled`, XenServer applies a filtering rule so that the VIF drops all traffic.

-`default-locking-mode=unlocked`, XenServer removes all the filtering rules associated with the VIF. By default, the default locking mode parameter is set to `unlocked`.

For information about the `default-locking-mode` parameter, see [Network commands](#).

The default locking mode of the network has no effect on attached VIFs whose locking state is anything other than `network_default`.

Note:

You cannot change the `default-locking-mode` of a network that has active VIFs attached to it.

- **Locked.** XenServer applies filtering rules so that only traffic sent to/from the specified MAC and IP addresses is allowed to be sent out through the VIF. In this mode, if no IP addresses are specified, the VM cannot send any traffic through that VIF, on that network.

To specify the IP addresses from which the VIF accepts traffic, use the IPv4 or IPv6 IP addresses by using the `ipv4_allowed` or `ipv6_allowed` parameters.

- **Unlocked.** All network traffic can pass through the VIF. That is, no filters are applied to any traffic going to or from the VIF.
- **Disabled.** No traffic is allowed to pass through the VIF. (That is, XenServer applies a filtering rule so that the VIF drops all traffic.)

Configure switch port locking This section provides three different procedures:

- Restrict VIFs to use a specific IP address
- Add an IP address to an existing restricted list. For example, to add an IP address to a VIF when the VM is running and connected to the network (for example, if you are taking a network offline temporarily).
- Remove an IP address from an existing restricted list

If a VIF's locking-mode is set to **locked**, it can only use the addresses specified in the **ipv4-allowed** or **ipv6-allowed** parameters.

Because, in some relatively rare cases, VIFs may have more than one IP address, it is possible to specify multiple IP addresses for a VIF.

You can perform these procedures before or after the VIF is plugged in (or the VM is started).

Change the default-locking mode to locked, if it is not using that mode already, by running the following command:

```
1 xe vif-param-set uuid=vif-uuid locking-mode=locked
```

The **vif-uuid** represents the UUID of the VIF you want to allow to send traffic. To obtain the UUID, run the **xe vif-list** command on the host. **vm-uuid** Indicates the virtual machine for which the information appears. The device ID indicates the device number of the VIF.

Run the **vif-param-set** command to specify the IP addresses from which the virtual machine can send traffic. Do one or more of the following:

- Specify one or more IPv4 IP addresses destinations. For example:

```
1 xe vif-param-set uuid=vif-uuid ipv4-allowed=comma separated list  
  of ipv4-addresses
```

- Specify one or more IPv6 IP addresses destinations. For example:

```
1 xe vif-param-set uuid=vif-uuid ipv6-allowed=comma separated list  
  of ipv6-addresses
```

You can specify multiple IP addresses by separating them with a comma, as shown in the preceding example.

After performing the procedure to restrict a VIF to using a specific IP address, you can add one or more IP addresses the VIF can use.

Run the `vif-param-add` command to add the IP addresses to the existing list. Do one or more of the following:

- Specify the IPv4 IP address. For example:

```
1 xe vif-param-add uuid=vif-uuid ipv4-allowed=comma separated list  
  of ipv4-addresses
```

- Specify the IPv6 IP address. For example:

```
1 xe vif-param-add uuid=vif-uuid ipv6-allowed=comma separated list  
  of ipv6-addresses
```

If you restrict a VIF to use two or more IP addresses, you can delete one of those IP addresses from the list.

Run the `vif-param-remove` command to delete the IP addresses from the existing list. Do one or more of the following:

- Specify the IPv4 IP address to delete. For example:

```
1 xe vif-param-remove uuid=vif-uuid ipv4-allowed=comma separated  
  list of ipv4-addresses
```

- Specify the IPv6 IP address to delete. For example:

```
1 xe vif-param-remove uuid=vif-uuid ipv6-allowed=comma separated  
  list of ipv6-addresses
```

Prevent a virtual machine from sending or receiving traffic from a specific network The following procedure prevents a virtual machine from communicating through a specific VIF. As a VIF connects to a specific XenServer network, you can use this procedure to prevent a virtual machine from sending or receiving any traffic from a specific network. This provides a more granular level of control than disabling an entire network.

If you use the CLI command, you do not need to unplug the VIF to set the VIF's locking mode. The command changes the filtering rules while the VIF is running. In this case, the network connection still appears to be present, however, the VIF drops any packets the VM attempts to send.

Tip:

To find the UUID of a VIF, run the `xe vif-list` command on the host. The device ID indicates the device number of the VIF.

To prevent a VIF from receiving traffic, disable the VIF connected to the network from which you want to stop the VM from receiving traffic:

```
1 xe vif-param-set uuid=vif-uuid locking-mode=disabled
```

You can also disable the VIF in XenCenter by selecting the virtual network interface in the VM's Networking tab and clicking Deactivate.

Remove a VIF's restriction to an IP address To revert to the default (original) locking mode state, use the following procedure. By default, when you create a VIF, XenServer configures it so that it is not restricted to using a specific IP address.

To revert a VIF to an unlocked state, change the VIF default-locking mode to unlocked. If it is not using that mode already, run the following command:

```
1 xe vif-param-set uuid=vif_uuid locking-mode=unlocked
```

Simplify VIF locking mode configuration in the Cloud Rather than running the VIF locking mode commands for each VIF, you can ensure all VIFs are disabled by default. To do so, you must change the packet filtering at the network level. Changing the packet filtering causes the XenServer network to determine how packets are filtered, as described in the previous section *How switch-port locking works*.


Specifically, a network's `default-locking-mode` setting determines how new VIFs with default settings behave. Whenever a VIF's `locking-mode` is set to `default`, the VIF refers to the network-locking mode (`default-locking-mode`) to determine if and how to filter packets traveling through the VIF:

- **Unlocked.** When the network `default-locking-mode` parameter is set to `unlocked`, XenServer lets the VM send traffic to any IP address on the network the VIF connects to.
- **Disabled.** When the `default-locking-mode` parameter is set to `disabled`, XenServer applies a filtering rule so that the VIF drops all traffic.

By default, the `default-locking-mode` for all networks created in XenCenter and using the CLI are set to `unlocked`.

By setting the VIF's locking mode to its default (`network_default`), you can create a basic default configuration (at the network level) for all newly created VIFs that connect to a specific network.

This illustration shows how, when a VIF's `locking-mode` is set to its default setting (`network_default`), the VIF uses the network `default-locking-mode` to determine its behavior.



For example, by default, VIFs are created with their `locking-mode` set to `network_default`. If you set a network's `default-locking-mode=disabled`, any new VIFs for which you have not configured the locking mode are disabled. The VIFs remain disabled until you either (a) change the individual VIF's `locking-mode` parameter or (b) explicitly set the VIF's `locking-mode` to 'unlocked'. This is helpful when you trust a specific VM enough so you do not want to filter its traffic at all.

To change a network's default locking mode setting:

After creating the network, change the default-locking mode by running the following command:

```
1 xe network-param-set uuid=network-uuid default-locking-mode=[unlocked | disabled]
```

Note:

To get the UUID for a network, run the `xe network-list` command. This command displays the UUIDs for all the networks on the host on which you ran the command.

To check a network's default locking mode setting:

Run one of the following commands:

```
1 xe network-param-get uuid=network-uuid param-name=default-locking-mode
```

OR

```
1 xe network-list uuid=network-uuid params=default-locking-mode
```

Use network settings for VIF traffic filtering The following procedure instructs a VIF on a virtual machine to use the XenServer network **default-locking-mode** settings on the network itself to determine how to filter traffic.

1. Change the VIF locking state to **network_default**, if it is not using that mode already, by running the following command:

```
1 xe vif-param-set uuid=vif_uuid locking-mode=network_default
```

2. Change the default-locking mode to **unlocked**, if it is not using that mode already, by running the following command:

```
1 xe network-param-set uuid=network-uuid default-locking-mode=unlocked
```

Dom0 network interface naming on XenServer 9

February 19, 2026

Overview

New in XenServer® 9: The control domain (dom0) adopts systemd predictable network interface names and consolidates interface handling in Toolstack networkd. In previous releases, multiple components renamed dom0 interfaces to the legacy ethN pattern. In XenServer 9, dom0 interface names are no longer forcibly changed to ethN.

This change improves stability and maintainability and aligns dom0 with industry-standard naming.

What's new in XenServer 9

- Predictable interface names in dom0 based on stable hardware attributes (for example, eno12429np0).
- Consolidated interface management in Toolstack networkd.
- Removal of forced renaming to ethN (for example, eth0, eth1).

Differences from XenServer 8.4

- XenServer 8.4: Several components cooperated to sort and rename dom0 interfaces to ethN.

- XenServer 9: Dom0 uses systemd predictable names and a single component manages interface naming and state.

Note:

Predictable names are derived from hardware and topology. Exact names vary by platform.

What you see in the product

- Interface names in dom0 follow systemd conventions. For example, eth0 might appear as eno12429np0. The device field of the XAPI PIF reflects the new name.
- The default network name-label for the pool-wide network that was associated with eth0 is now Pool-wide network 0.

Benefits

- Consolidated interface management improves product stability and maintainability.
- Predictable names reduce ambiguity across reboots and hardware changes.

Migrate your automation and scripts

When you upgrade to XenServer 9, scripts or tools that assume ethN device names might fail. Update automation to discover interfaces dynamically instead of hard-coding ethN.

Examples:

- List dom0 interfaces by using Linux tools:

```
1 ip -o link show
```

- Discover PIFs and devices by using xe:

```
1 xe pif-list host-uuid=<host-uuid> params=uuid,device,MAC,management
2 xe network-list params=uuid,name-label
```

Important:

Do not rely on ethN names in documentation or scripts. Dom0 interface names are not guaranteed to use the ethN pattern.

Troubleshoot

- A script fails after upgrade because it references eth0 or eth1.

- Update the script to discover the PIF by MAC address, management flag, or network name-label.
- You cannot find the management interface.
 - In dom0, run `ip -o link show` to list devices. In XenCenter, check Host Properties > NICs. In xe, run `xe pif-list management=true params=uuid,device,MAC`.

Collect logs and contact Support if issues persist.

Troubleshoot networking

June 22, 2026

If you are experiencing problems with configuring networking, first ensure that you have not directly changed any of the control domain `ifcfg-*` files. The control domain host agent manages the `ifcfg` files directly, and any changes are overwritten.

Diagnosing network corruption

Some network card models require firmware upgrades from the vendor to work reliably under load, or when certain optimizations are turned on. If you see corrupted traffic to VMs, try to obtain the latest firmware from your vendor and then use it to update your hardware.

If the problem still persists, then you can use the CLI to disable receive or transmit offload optimizations on the physical interface.

Warning:

Disabling receive or transmit offload optimizations can result in a performance loss and increased CPU usage.

First, determine the UUID of the physical interface. You can filter on the `device` field as follows:

```
1 xe pif-list device=<network-interface>
```

Next, set the following parameter on the PIF to disable TX offload:

```
1 xe pif-param-set uuid=pif_uuid other-config:ethtool-tx=off
```

Finally, replug the PIF or restart the host for the change to take effect.

Emergency network reset

Incorrect networking settings can cause loss of network connectivity. When there is no network connectivity, XenServer host can become inaccessible through XenCenter® or remote SSH. Emergency Network Reset provides a simple mechanism to recover and reset a host's networking.

The Emergency network reset feature is available from the CLI using the `xe-reset-networking` command, and within the **Network and Management Interface** section of `xsconsole`.

Incorrect settings that cause a loss of network connectivity include renaming network interfaces, creating bonds or VLANs, or mistakes when changing the management interface. For example, typing the wrong IP address. You may also want to run this utility in the following scenarios:

- When a rolling pool upgrade, manual upgrade, hotfix installation, or driver installation causes a lack of network connectivity, or
- If a pool coordinator or host in a resource pool is unable to contact with other hosts.

Use the `xe-reset-networking` utility only in an emergency because it deletes the configuration for all PIFs, bonds, VLANs, and tunnels associated with the host. Guest Networks and VIFs are preserved. As part of this utility, VMs are shut down forcefully. Before running this command, cleanly shut down the VMs where possible. Before you apply a reset, you can change the management interface and specify which IP configuration, DHCP, or Static can be used.

If the pool coordinator requires a network reset, reset the network on the pool coordinator first before applying a network reset on pool members. Apply the network reset on all remaining hosts in the pool to ensure that the pool's networking configuration is homogeneous. Network homogeneity is an important factor for live migration.

Note:

If the pool coordinator's IP address (the management interface) changes as a result of a network reset or `xe host-management-reconfigure`, apply the network reset command to other hosts in the pool. This is to ensure that the pool members can reconnect to the pool coordinator on its new IP address. In this situation, the IP address of the pool coordinator must be specified.

Network reset is NOT supported when High Availability is enabled. To reset network configuration in this scenario, you must first manually disable high availability, and then run the network reset command.

Verifying the network reset

After you specify the configuration mode to be used after the network reset, `xsconsole` and the CLI display settings that will be applied after host reboot. It is a final chance to modify before applying the emergency network reset command. After restart, the new network configuration can be verified in

XenCenter and `xsconsole`. In XenCenter, with the host selected, select the **Networking** tab to see the new network configuration. The Network and Management Interface section in `xsconsole` display this information.

Note:

Run emergency network reset on other pool members to replicate bonds, VLANs, or tunnels from the pool coordinator's new configuration.

Using the CLI for network reset

The following table shows the available optional parameters which can be used by running the `xe-reset-networking` command.

Warning:

Users are responsible to ensure the validity of parameters for the `xe-reset-networking` command, and to check the parameters carefully. If you specify invalid parameters, network connectivity and configuration can be lost. In this situation, we advise that you rerun the command `xe-reset-networking` without using any parameters.

Resetting the networking configuration of a whole pool **must** begin on the pool coordinator, followed by network reset on all remaining hosts in the pool.

Parameter	Required/Optional	Description
<code>-m, --master</code>	Optional	IP address of the pool coordinator's management interface. Defaults to the last known pool coordinator's IP address.
<code>--device</code>	Optional	Device name of the management interface. Defaults to the device name specified during installation.
<code>--mode=static</code>	Optional	Enables the following four networking parameters for static IP configuration for the management interface. If not specified, networking is configured using DHCP.

Parameter	Required/Optional	Description
<code>--ip</code>	Required, if <code>mode=static</code>	IP address for the host's management interface. Only valid if <code>mode=static</code> .
<code>--netmask</code>	Required, if <code>mode=static</code>	Netmask for the management interface. Only valid if <code>mode=static</code> .
<code>--gateway</code>	Optional	Gateway for the management interface. Only valid if <code>mode=static</code> .
<code>--dns</code>	Optional	DNS Server for the management interface. Only valid if <code>mode=static</code> .
<code>--vlan</code>	Optional	VLAN tag for the management interface. Defaults to the VLAN tag specified during installation.

Pool coordinator command-line examples Examples of commands that can be applied on a pool coordinator:

To reset networking for DHCP configuration:

```
1 xe-reset-networking
```

To reset networking for Static IP configuration:

```
1 xe-reset-networking --mode= static --ip=ip-address \
2   --netmask=netmask --gateway=gateway \
3   --dns=dns
```

To reset networking for DHCP configuration if another interface became the management interface after initial setup:

```
1 xe-reset-networking --device=device-name
```

To reset networking for Static IP configuration if another interface became the management interface after initial setup:

```
1 xe-reset-networking --device=device-name --mode=static \
2   --ip=ip-address --netmask=netmask \
3   --gateway=gateway --dns=dns
```

To reset networking for management interface on VLAN:

```
1 xe-reset-networking --vlan=VLAN TAG
```

Note:

The `reset-network` command can also be used along with the IP configuration settings.

Pool member command-line examples All previous examples also apply to pool members. Additionally, the pool coordinator's IP address can be specified (which is necessary if it has changed.)

To reset networking for DHCP configuration:

```
1 xe-reset-networking
```

To reset networking for DHCP if the pool coordinator's IP address was changed:

```
1 xe-reset-networking --master=pool-coordinator-ip-address
```

To reset networking for Static IP configuration, assuming the pool coordinator's IP address didn't change:

```
1 xe-reset-networking --mode=static --ip=ip-address --netmask=netmask \  
2 --gateway=gateway --dns=dns
```

To reset networking for DHCP configuration if the management interface and the pool coordinator's IP address was changed after initial setup:

```
1 xe-reset-networking --device=device-name --master=pool-coordinator-ip-  
address
```


Storage

June 30, 2026

This article describes how physical storage hardware maps to virtual machines (VMs), and the software objects used by XenServer® to perform storage-related tasks. It also provides information about the storage types available to enable you to make the right choice for the storage hardware to use in your environment.

Storage concepts

The following image shows how the storage objects described in this section are related:



Storage repositories (SRs)

A Storage Repository (SR) is a particular storage target, in which Virtual Machine (VM) Virtual Disk Images (VDIs) are stored. A VDI is a storage abstraction that represents a virtual hard disk drive (HDD). To learn about the types of SRs supported by XenServer, see [Storage repository types](#).

The SR and VDI abstractions allow for advanced storage features to be exposed on storage targets that support them. For example, advanced features such as *thin provisioning*, VDI snapshots, and fast cloning. For storage subsystems that don't support advanced operations directly, a software stack that implements these features is provided.

A storage repository is a persistent, on-disk data structure. For SR types that use an underlying block device, the process of creating an SR involves erasing any existing data on the specified storage target. Other storage types such as NFS, create a container on the storage array in parallel to existing SRs.

Each XenServer host can use multiple SRs and different SR types simultaneously. These SRs can be shared between hosts or dedicated to particular hosts. Shared storage is pooled between multiple hosts within a defined resource pool. A shared SR must be network accessible to each host in the pool. Shared storage cannot be shared between multiple pools.

SR commands provide operations for creating, destroying, resizing, cloning, connecting and discovering the individual VDIs that they contain. CLI operations to manage storage repositories are described in [SR commands](#).

Virtual disk image (VDI)

A virtual disk image (VDI) is a storage abstraction that represents a virtual hard disk drive (HDD). VDIs are the fundamental unit of virtualized storage in XenServer. VDIs are persistent, on-disk objects that

exist independently of XenServer hosts. For more information about the data formats supported for virtual disks on XenServer, see [Virtual disk data formats](#).

CLI operations to manage VDIs are described in [VDI commands](#). The on-disk representation of the data differs by SR type. A separate storage plug-in interface for each SR, called the SM API, manages the data.

Physical block devices (PBDs)

Physical block devices represent the interface between a physical server and an attached SR. PBDs are connector objects that allow a given SR to be mapped to a host. PBDs store the device configuration fields that are used to connect to and interact with a given storage target. For example, NFS device configuration includes the IP address of the NFS server and the associated path that the XenServer host mounts. PBD objects manage the run-time attachment of a given SR to a given XenServer host. CLI operations relating to PBDs are described in [PBD commands](#).

Virtual block devices (VBDs)

Virtual Block Devices are connector objects (similar to the PBD described above) that allows mappings between VDIs and VMs. In addition to providing a mechanism for attaching a VDI into a VM, VBDs allow for the fine-tuning of parameters regarding the disk I/O priority and statistics of a given VDI, and whether that VDI can be booted. CLI operations relating to VBDs are described in [VBD commands](#).

Storage repository types

XenServer supports a range of storage types for both locally and remotely connected SRs.

Warning:

XenServer does not support snapshots at the external SAN-level of a LUN for any SR type.

Local SRs

A local storage repository (SR) is connected to a single host and is not shared between all hosts in a pool. The local physical storage hardware can be a hard disk drive (HDD) or a solid state drive (SSD). It can connect to the host using any of the following methods: SATA, SCSI, SAS, NVMe.

XenServer hosts support the following types of local storage:

- [Local LVM](#)
- [Local EXT3/EXT4](#)
- [Local XFS](#)

Remote SRs

XenServer supports remotely connected storage using the following drives: iSCSI, NFS, SAS, SMB (version 3 only), Fibre Channel.

Note:

NVMe over Fibre Channel and NVMe over TCP are not supported.

The following remotely connected storage types can be used to create shared SRs for your XenServer pool.

ISO libraries:

- [ISO library on NFS storage](#)
- [ISO library on CIFS/SMB storage](#)

File-based storage:

- [NFS storage](#)
- [SMB storage](#)

Block-based storage:

- [Software iSCSI storage](#)
- [Hardware HBA storage](#)
- [Thin-provisioned shared GFS2 block storage](#)

Comparison of shared storage types

The following table compares the features of the different storage types supported for shared SRs. Consider these capabilities when choosing your storage solution.

	File-based (SMB/NFS)	Block-based (iSCSI/HBA) LVM	Block-based (iSCSI/HBA) GFS2
Virtual disk type	VHD	VHD	QCOW2
Maximum virtual disks per SR	20000	1000	20000
Maximum virtual disk size	2,040 GiB	2,040 GiB	16 TiB
Maximum pool size	32/64	32	16

Supported features

	File-based (SMB/NFS)	Block-based (iSCSI/HBA) LVM	Block-based (iSCSI/HBA) GFS2
Storage migration	Yes	Yes	No
Intellicache	Yes	Yes	No
Read caching	Yes	No	Yes
Thin provisioning (see note)	Yes	No	Yes
Disaster recovery	No	Yes	No

Note:

For block-based LVM storage, you can implement thin provisioning in the underlying SAN. However, this is not recommended with GFS2 SRs.

Virtual disk data formats

In general, there are the following types of mapping of physical storage to a VDI:

1. *Logical volume-based VHD on a LUN:* The default XenServer block-based storage inserts a logical volume manager on a disk. This disk is either a locally attached device (LVM) or a SAN attached LUN over either Fibre Channel, iSCSI, or SAS. VDIs are represented as volumes within the volume manager and stored in VHD format to allow thin provisioning of reference nodes on snapshot and clone.
2. *File-based QCOW2 on a LUN:* VM images are stored as thin-provisioned QCOW2 format files on a GFS2 shared-disk filesystem on a LUN attached over either iSCSI software initiator or Hardware HBA.
3. *File-based VHD on a filesystem:* VM images are stored as thin-provisioned VHD format files on either a local non-shared filesystem (EXT3/EXT4 type SR), a shared NFS target (NFS type SR), or a remote SMB target (SMB type SR).
4. *File-based QCOW2 on a filesystem:* VM images are stored as thin-provisioned QCOW2 format files on a local non-shared XFS filesystem.

VDI types

For GFS2 and XFS SRs, QCOW2 VDIs are created.

For other SR types, VHD-format VDIs are created. You can opt to use raw at the time you create the VDI. This option can only be specified by using the xe CLI.

Note:

If you create a raw VDI on an LVM-based SR or HBA/LUN-per-VDI SR, it might allow the owning VM to access data that was part of a previously deleted VDI (of any format) belonging to any VM. We recommend that you consider your security requirements before using this option.

Raw VDIs on an NFS, EXT, or SMB SR do not allow access to the data of previously deleted VDIs belonging to any VM.

To check if a VDI was created with `type=raw`, check its `sm-config` map. The `sr-param-list` and `vd-param-list` `xe` commands can be used respectively for this purpose.

Create a raw virtual disk by using the `xe` CLI

1. Run the following command to create a VDI given the UUID of the SR you want to place the virtual disk in:

```
1 xe vdi-create sr-uuid=sr-uuid type=user virtual-size=virtual-size
2 \ name-label=VDI name sm-config:type=raw
```

2. Attach the new virtual disk to a VM. Use the disk tools within the VM to partition and format, or otherwise use the new disk. You can use the `vbd-create` command to create a VBD to map the virtual disk into your VM.

Convert between VDI formats

It is not possible to do a direct conversion between the raw and VHD formats. Instead, you can create a VDI (either raw, as described above, or VHD) and then copy data into it from an existing volume. Use the `xe` CLI to ensure that the new VDI has a virtual size at least as large as the VDI you are copying from. You can do this by checking its `virtual-size` field, for example by using the `vd-param-list` command. You can then attach this new VDI to a VM and use your preferred tool within the VM to do a direct block-copy of the data. For example, standard disk management tools in Windows or the `dd` command in Linux. If the new volume is a VHD volume, use a tool that can avoid writing empty sectors to the disk. This action can ensure that space is used optimally in the underlying storage repository. A file-based copy approach may be more suitable.

VHD-based and QCOW2-based VDIs

VHD and QCOW2 images can be *chained*, allowing two VDIs to share common data. In cases where a VHD-backed or QCOW2-backed VM is cloned, the resulting VMs share the common on-disk data at

the time of cloning. Each VM proceeds to make its own changes in an isolated copy-on-write version of the VDI. This feature allows such VMs to be quickly cloned from templates, facilitating very fast provisioning and deployment of new VMs.

As VMs and their associated VDIs get cloned over time this creates trees of chained VDIs. When one of the VDIs in a chain is deleted, XenServer rationalizes the other VDIs in the chain to remove unnecessary VDIs. This *coalescing* process runs asynchronously. The amount of disk space reclaimed and time taken to perform the process depends on the size of the VDI and the amount of shared data. For more information, see Notes on coalesce.

Both the VHD and QCOW2 formats support *thin provisioning*. The image file is automatically extended in fine granular chunks as the VM writes data into the disk. For file-based VHD and GFS2-based QCOW2, this approach has the considerable benefit that VM image files take up only as much space on the physical storage as required. With LVM-based VHD, the underlying logical volume container must be sized to the virtual size of the VDI. However unused space on the underlying copy-on-write instance disk is reclaimed when a snapshot or clone occurs. The difference between the two behaviors can be described in the following way:

- For *LVM-based VHD images*, the difference disk nodes within the chain consume only as much data as has been written to disk. However, the leaf nodes (VDI clones) remain fully inflated to the virtual size of the disk. Snapshot leaf nodes (VDI snapshots) remain deflated when not in use and can be attached Read-only to preserve the deflated allocation. Snapshot nodes that are attached in read-write mode are fully inflated on attach, and deflated on detach.
- For *file-based VHDs and GFS2-based QCOW2 images*, all nodes consume only as much data as has been written. The leaf node files grow to accommodate data as it is actively written. If a 100 GB VDI is allocated for a VM and an OS is installed, the VDI file is physically only the size of the OS data on the disk, plus some minor metadata overhead.

When cloning VMs based on a single VHD or QCOW2 template, each child VM forms a chain where new changes are written to the new VM. Old blocks are directly read from the parent template. If the new VM was converted into a further template and more VMs cloned, then the resulting chain results in degraded performance. XenServer supports a maximum chain length of 30. Do not approach this limit without good reason. If in doubt, “copy” the VM using XenCenter or use the `vm-copy` command, which resets the chain length back to 0.

Notes on coalesce Only one coalescing process is ever active for an SR. This process thread runs on the SR pool coordinator. While the coalescing process is active, extra space in the SR might be required to accommodate data being merged from one level of the delta tree to another. After this process, the redundant level is removed and space is freed.

If you have critical VMs running on the pool coordinator, you can take the following steps to mitigate against occasional slow I/O:

- Migrate the VM to a host other than the SR pool coordinator
- Set the disk I/O priority to a higher level, and adjust the scheduler. For more information, see [Virtual disk I/O request prioritization](#).

Create a storage repository

February 18, 2026

You can create local and remote storage repositories by using the `xe` CLI. Here are two basic steps to create a storage repository for use on a host by using the CLI:

1. Determine values for any required parameters by probing the SR type with the command `xe sr-probe`.
2. Create the SR with the command `xe sr-create`.

This command creates an SR on the storage substrate (potentially destroying any existing data). It also creates the SR API object and a corresponding PBD record, enabling VMs to use the storage. On successful creation of the SR, the PBD is automatically plugged. If the flag `shared=true` is set, a PBD record is created and plugged for every XenServer host in the resource pool.

These steps differ in detail depending on the type of SR being created. In all examples, the `sr-create` command returns the UUID of the created SR if successful.

Alternatively, you can use the **New Storage Repository** wizard in XenCenter to create remote storage repositories (SRs). The wizard guides you through the configuration steps. For more information, see the [XenCenter documentation](#).

For examples of how to create an SR, see the article for the storage type you use:

- [Local storage](#)

ISO libraries:

- [ISO library on NFS storage](#)
- [ISO library on CIFS/SMB storage](#)

File-based storage:

- [NFS storage](#)
- [SMB storage](#)

Block-based storage:

- [Software iSCSI storage](#)

- [Hardware HBA storage](#)
- [Thin-provisioned shared GFS2 block storage](#)

All XenServer SR types support VDI resize, fast cloning, and snapshot. SRs based on the LVM SR type (local, iSCSI, or HBA) provide thin provisioning for snapshot and hidden parent nodes, but not for other disks, such as MCS per-VM delta disks. The other SR types (EXT3/EXT4, NFS, GFS2, and XFS) support full thin provisioning, including for virtual disks that are active.

Warnings:

- When VHD VDIs are not attached to a VM, for example for a VDI snapshot, they are stored as thinly provisioned by default. If you attempt to reattach the VDI, ensure that there is sufficient disk-space available for the VDI to become thickly provisioned. VDI clones are thickly provisioned.
- XenServer® does not support snapshots at the external SAN-level of a LUN for any SR type.
- Do not attempt to create an SR where the LUN ID of the destination LUN is greater than 255. Ensure that your target exposes the LUN with a LUN ID that is less than or equal to 255 before using this LUN to create an SR.
- If you use thin provisioning on a file-based SR, ensure that you monitor the free space on your SR. If the SR usage grows to 100%, further writes from VMs fail. These failed writes can cause the VM to freeze or crash, and can lead to data corruption.

Probe an SR

The `sr-probe` command can be used in the following ways:

- To identify unknown parameters for use in creating an SR
- To return a list of existing SRs

In both cases `sr-probe` works by specifying an SR type and one or more `device-config` parameters for that SR type. If an incomplete set of parameters is supplied, the `sr-probe` command returns an error message indicating parameters are missing and the possible options for the missing parameters. When a complete set of parameters is supplied, a list of existing SRs is returned. All `sr-probe` output is returned as XML.

For example, a known iSCSI target can be probed by specifying its name or IP address. The set of IQNs available on the target is returned:

```
1      xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10
2
3      Error code: SR_BACKEND_FAILURE_96
4      Error parameters: , The request is missing or has an incorrect
      target IQN parameter, \
```



```

5      <?xml version="1.0" ?>
6      <iscsi-target-iqns>
7          <TGT>
8              <Index>
9                  0
10             </Index>
11             <IPAddress>
12                 192.168.1.10
13             </IPAddress>
14             <TargetIQN>
15                 iqn.192.168.1.10:filer1
16             </TargetIQN>
17         </TGT>
18     </iscsi-target-iqns>

```

Probing the same target again and specifying both the name/IP address and desired IQN returns the set of **SCSIids** (LUNs) available on the target/IQN.

```

1      xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10 \
2      device-config:targetIQN=iqn.192.168.1.10:filer1
3
4      Error code: SR_BACKEND_FAILURE_107
5      Error parameters: , The SCSIid parameter is missing or incorrect, \
6      <?xml version="1.0" ?>
7      <iscsi-target>
8          <LUN>
9              <vendor>
10                 IET
11             </vendor>
12             <LUNid>
13                 0
14             </LUNid>
15             <size>
16                 42949672960
17             </size>
18             <SCSIid>
19                 1494554000000000000000000000000020000000b70200000f000000
20             </SCSIid>
21         </LUN>
22     </iscsi-target>

```

Probing the same target and supplying all three parameters returns a list of SRs that exist on the LUN, if any.

```

1      xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10 \
2      device-config:targetIQN=192.168.1.10:filer1 \
3      device-config:SCSIid=1494554000000000000000000000000020000000
4      b70200000f000000
5
6      <?xml version="1.0" ?>
7      <SRlist>
8          <SR>
9              <UUID>

```

```

9          3f6e1ebd-8687-0315-f9d3-b02ab3adc4a6
10         </UUID>
11         <Devlist>
12             /dev/disk/by-id/scsi-14945540000000000000000002000000
              b7020000f000000
13         </Devlist>
14     </SR>
15 </SRlist>

```

The following parameters can be probed for each SR type:

SR type	The <code>device-config</code> parameters, in order of dependency	Can be probed?	Required for <code>sr-create</code> ?
lvmoiscsi	target	No	Yes
	chapuser	No	No
	chappassword	No	No
	targetIQN	Yes	Yes
	SCSIid	Yes	Yes
lvmohba	SCSIid	Yes	Yes
lvmofcoe	SCSIid	Yes	Yes
nfs	server	No	Yes
	serverpath	Yes	Yes
smb	server	No	Yes
	username	No	No
	password	No	No
lvm	device	No	Yes
ext	device	No	Yes

For information about probing a GFS2 SR, see [Create a GFS2 SR](#).

Local storage

February 18, 2026

A local storage repository (SR) is connected to a single host and is not shared between all hosts in a pool. The local physical storage hardware can be a hard disk drive (HDD) or a solid state drive (SSD) and connect to the host using any of the following methods: SATA, SCSI, SAS, NVMe.

XenServer® hosts support the following types of local storage:

- Local LVM
- Local EXT3/EXT4
- Local XFS

XenServer also uses the udev type to represent devices plugged in as VDIs. For more information, see [udev](#).

Note:

Local SRs of type `lvm`, `ext`, and `xfs` can only be created using the `xe` CLI. After creation, you can manage all SR types by either XenCenter or the `xe` CLI.

Local LVM

The Local LVM type presents disks within a locally attached Volume Group. We recommend that you only attach one local SR per host. The maximum VDI size for this SR is 2,040 GiB.

By default, XenServer uses the local disk on the physical host on which it is installed. The Linux Logical Volume Manager (LVM) is used to manage VM storage. A VDI is implemented in VHD format in an LVM logical volume of the specified size.

Note:

The block size of an LVM LUN must be 512 bytes. To use storage with 4 KB physical blocks, the storage must also support emulation of 512-byte allocation blocks (the logical block size must be 512 bytes).

LVM performance considerations

The snapshot and fast clone functionality for LVM-based SRs comes with an inherent performance overhead. When optimal performance is required, XenServer supports the creation of VDIs in the *raw* format in addition to the default VHD format. The XenServer snapshot functionality is not supported on raw VDIs.

Warning:

Do not try to snapshot a VM that has `type=raw` disks attached. This action can result in a partial

snapshot being created. In this situation, you can identify the orphan snapshot VDIs by checking the `snapshot-of` field and then deleting them.

Creating a local LVM SR

An LVM SR is created by default on host install.

Device-config parameters for LVM SRs are:

Parameter Name	Description	Required?
<code>device</code>	Device name on the local host to use for the SR. You can also provide a comma-separated list of names.	Yes

To create a local LVM SR on `/dev/disk/<id>`, use the following command.

```
1  xe sr-create host-uuid=valid_uuid content-type=user \  
2  name-label="Example Local LVM SR" shared=false \  
3  device-config:device=/dev/disk/<id> type=lvm
```

Local EXT3/EXT4

Using EXT3/EXT4 enables thin provisioning on local storage. However, the default storage repository type is LVM as it gives a consistent write performance and prevents storage over-commit. The maximum VDI size for this SR is 2,040 GiB.

If you use EXT3/EXT4, you might see reduced performance in the following cases:

- When carrying out VM lifecycle operations such as VM create and suspend/resume
- When creating large files from within the VM

Local disk EXT3/EXT4 SRs must be configured using the XenServer CLI.

Whether a local EXT SR uses EXT3 or EXT4 depends on what version of XenServer created it:

- If you created the local EXT SR on a version of Citrix Hypervisor or XenServer prior to XenServer 8.4, and then upgraded, it uses EXT3.
- If you created the local EXT SR on XenServer 8.4 or later, it uses EXT4.

Note:
The block size of an EXT3/EXT4 disk must be 512 bytes. To use storage with 4 KB physical blocks, the storage must also support emulation of 512-byte allocation blocks (the logical block size

must be 512 bytes).

Creating a local EXT4 SR (ext)

Device-config parameters for EXT SRs:

Parameter Name	Description	Required?
device	Device name on the local host to use for the SR. You can also provide a comma-separated list of names.	Yes

To create a local EXT4 SR on `/dev/disk/<id>`, use the following command:

```
1 xe sr-create host-uuid=valid_uuid content-type=user \  
2   name-label="Example Local EXT4 SR" shared=false \  
3   device-config:device=/dev/disk/<id> type=ext
```

Local XFS

Using XFS enables thin provisioning on local storage. The local XFS type allows you to create local storage devices with 4 KB physical blocks without requiring a logical block size of 512 bytes. The maximum supported VDI size for this SR is 16 TiB.

Constraints

XFS SRs have the following constraints:

- Live Storage migration cannot be used on VMs whose VDIs are on an XFS SR. However, you can manually copy VDIs between SRs.
- XFS SRs cannot be used as the local cache storage for IntelliCache™.
- Trim/unmap is not supported on XFS SRs.
- You cannot export VDIs that are greater than 2 TiB as VHD or OVA/OVF. However, you can export VMs with VDIs larger than 2 TiB in XVA format.

Creating a local XFS SR

Device-config parameters for XFS SRs:

Parameter Name	Description	Required?
<code>device</code>	Device name on the local host to use for the SR. You can also provide a comma-separated list of names.	Yes

To create a local XFS SR on `/dev/disk/<id>`, use the following command:

```
1  xe sr-create host-uuid=valid_uuid content-type=user \  
2  name-label="Example Local XFS SR" shared=false \  
3  device-config:device=/dev/disk/<id> type=xfs
```

udev

The `udev` type represents devices plugged in using the `udev` device manager as VDIs. XenServer has two SRs of type `udev` that represent removable storage. One is for the CD or DVD disk in the physical CD or DVD-ROM drive of the XenServer host. The other is for a USB device plugged into a USB port of the XenServer host. VDIs that represent the media come and go as disks or USB sticks are inserted and removed.

ISO library on NFS

December 17, 2025

The `ISO` type handles CD images stored as files in `ISO` format. This storage repository (SR) type is useful for creating shared `ISO` libraries.

An `ISO` SR of type `nfs_iso` handles CD images stored as files in `ISO` format available as an `NFS` share.

If you do not specify the storage type to use for the SR, XenServer uses the `location` device config parameter to decide the type.

For more information about using the `NFS` SR types, see [NFS storage](#).

If you are creating an SR for IP-based storage (iSCSI or NFS), you can configure one of the following as the storage network: the NIC that handles the management traffic or a new NIC for the storage traffic. To assign an IP address to a NIC, see [Configure a dedicated storage NIC](#).

Device-config parameters for ISO SRs on NFS:

Parameter Name	Description	Required?
<code>location</code>	Path to the mount.	Yes
<code>type</code>	Storage type to use for the SR. For NFS, use <code>nfs_iso</code> .	No
<code>nfsversion</code>	For the storage type NFS, the version of the NFS protocol to use: 3, 4, 4.0, or 4.1.	No

Create an ISO SR on NFS

xe CLI

For storage repositories that store a library of ISOs, the `content-type` parameter must be set to `iso`, for example:

```
1 xe sr-create host-uuid=valid_uuid content-type=iso type=iso name-
   label="Example ISO SR" \
2 device-config:location=<path_to_mount> device-config:type=nfs_iso
```

In XenCenter®

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.
2. Under **ISO library**, select **NFS ISO** as the storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box.

Click **Next** to continue.

4. On the **Location** page, specify the ISO storage target details:
 - **Share Name:** For example, `server:/path` where `server` is the DNS name or IP address of the server computer, and `path` is a folder or file name.
 - **NFS Version:** Select the NFS version that the SR uses.

Note:

If the underlying storage array does not support NFSv4, NFSv3 is used to mount the share.

- 5. Click **Finish** to complete the new SR configuration and close the wizard.

ISO library on SMB/CIFS

December 17, 2025

The ISO type handles CD images stored as files in ISO format. This storage repository (SR) type is useful for creating shared ISO libraries.

The ISO SR type `cifs` uses the Windows File Sharing (SMB/CIFS) SR type to handle CD images stored as files in ISO format available as a Windows (SMB/CIFS) share.

If you do not specify the storage type to use for the SR, XenServer uses the `location` device config parameter to decide the type.

Device-config parameters for ISO SRs on CIFS/SMB:

Parameter Name	Description	Required?
<code>location</code>	Path to the mount.	Yes
<code>type</code>	Storage type to use for the SR: For CIFS/SMB, set to <code>cifs</code> .	No
<code>vers</code>	For the storage type CIFS/SMB, the version of SMB to use: 1.0 or 3.0. The default is 3.0.	No
<code>username</code>	For the storage type CIFS/SMB, if a user name is required for the Windows file server.	No
<code>cifspassword_secret</code>	(Recommended) For the storage type CIFS/SMB, you can pass a secret instead of a password for the Windows file server.	No

Parameter Name	Description	Required?
<code>cifspassword</code>	For the storage type CIFS/SMB, if a password is required for the Windows file server. We recommend that you use the <code>cifspassword_secret</code> parameter instead.	No

Note:

When running the `sr-create` command, we recommend that you use the `device-config:cifspassword_secret` argument instead of specifying the password on the command line. For more information, see [Secrets](#).

For more information about using the SMB SR type, see [SMB storage](#).

Create an ISO SR on CIFS/SMB**xe CLI**

We recommend that you use SMB version 3 to mount ISO SR on a Windows file server. Version 3 is selected by default because it is more secure and robust than SMB version 1.0.

Create an SMB version 3 ISO SR (using the recommended `cifspassword_secret` option) as

```

1  xe secret-create description="SMB ISO SR" value=<password>
2  xe sr-create content-type=iso type=iso shared=true device-config:
   location=<path_to_mount> \
3  device-config:username=<username> device-config:
   cifspassword_secret=<uuid returned from secret-create> \
4  device-config:type=cifs name-label="Example ISO SR"
```

If SMB version 1 is required due to filer limitations, you can mount ISO SR using SMB version 1 using the following command:

```

1  xe sr-create content-type=iso type=iso shared=true device-config:
   location=<path_to_mount> \
2  device-config:username=<username> device-config:cifspassword=<
   password> \
3  device-config:type=cifs device-config:vers=1.0 name-label="Example
   ISO SR"
```

In XenCenter®

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.
2. Under **ISO library**, select **Windows File Sharing (SMB/CIFS)** as the storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box.

Click **Next** to continue.

4. On the **Location** page, specify the ISO storage target details:
 - **Share Name:** For example, `\\server\sharename` where `server` is the DNS name or IP address of the server computer, and `sharename` is a folder or file name.
 - **Use different user name:** If you want to connect to an SMB server using a different user name, select this check box and then enter your login user name and password.
5. Click **Finish** to complete the new SR configuration and close the wizard.

NFS SRs

February 18, 2026

Shares on NFS servers (that support any version of NFSv4 or NFSv3) can be used immediately as an SR for virtual disks. VDIs are stored in the Microsoft VHD format only. Also, as these SRs can be shared, VDIs stored on shared SRs allow:

- VMs to be started on any XenServer® hosts in a resource pool
- VMs migrate between XenServer hosts in a resource pool using live migration (without noticeable downtime)

Important:

- For NFSv4, only the authentication type `AUTH_SYS` is supported.
- It is highly recommended for both NFS storage that a dedicated storage network be used, using at least two bonded links, ideally to independent network switches with redundant power supplies.
- All hosts in the Resource Pool must connect to the same filer IP in order to guarantee a consistent state being observed by all hosts. This is particularly important when High Avail-

ability is configured as inconsistencies relating to access of the High Availability statefile can result in hosts being fenced. See [High Availability](#).

- Care must be taken when using filers with multiple controllers, typically controller fail-over does not preserve the state of open files on the NFS share and thus a fail-over can cause any VMs running on the Storage Repository to crash or become read-only as they may lose access to their backing storage.
 - When performing updates to the filer controllers that may require a controller fail-over the NFS Storage Repository should be detached from the Resource Pool before initiating any changeover of the controllers.
 - After an unplanned fail-over it may be necessary to reboot VMs that have become unresponsive.

VDIs stored on file-based SRs are *thinly provisioned*. The image file is allocated as the VM writes data into the disk. This approach has the considerable benefit that the VM image files take up only as much space on the storage as is required. For example, if a 100 GB VDI is allocated for a VM and an OS is installed, the VDI file only reflects the size of the OS data written to the disk rather than the entire 100 GB.

VHD files can also be chained, allowing two VDIs to share common data. In cases where a file-based VM is cloned, the resulting VMs share the common on-disk data at the time of cloning. Each VM proceeds to make its own changes in an isolated copy-on-write version of the VDI. This feature allows file-based VMs to be quickly cloned from templates, facilitating fast provisioning and deployment of new VMs. The maximum supported length of VHD chains is 30.

File-based SRs and VHD implementations in XenServer assume that they have full control over the SR directory on the file server. Administrators must not modify the contents of the SR directory, as this action can risk corrupting the contents of VDIs.

XenServer has been tuned for enterprise-class storage that uses non-volatile RAM to provide fast acknowledgments of write requests while maintaining a high degree of data protection from failure.

Warning:

As VDIs on file-based SRs are created as thin-provisioned VDIs, administrators must ensure that the file-based SRs have enough disk space for all required VDIs. XenServer hosts do not enforce that the space required for VDIs on file-based SRs is present.

Ensure that you monitor the free space on your SR. If the SR usage grows to 100%, further writes from VMs fail. These failed writes can cause the VM to freeze or crash.

The following `device-config` parameters are used with NFS SRs:

Parameter Name	Description	Required?
<code>server</code>	IP address or host name of the NFS server	Yes
<code>serverpath</code>	Path, including the NFS mount point, to the NFS server that hosts the SR	Yes
<code>nfsversion</code>	Specifies the version of NFS to use. If you specify <code>nfsversion="4"</code> , the SR uses NFS v4.0, v4.1 or v4.2, depending on what is available. If you want to select a more specific version of NFS, you can specify <code>nfsversion="4.0"</code> and so on. Only one value can be specified for <code>nfsversion</code> .	No
<code>useUDP</code>	Configure the SR to use UDP rather than the default TCP.	No

Create a shared NFS SR (NFS)

Note:

If you attempt to attach a read-only NFS SR, this action fails with the following error message: "SR_BACKEND_FAILURE_461 - The file system for SR cannot be written to."

To create an NFS SR, you must provide the host name or IP address of the NFS server. You can create the SR on any valid destination path; use the `sr-probe` command to display a list of valid destination paths exported by the server.

In scenarios where XenServer is used with lower-end storage, it cautiously waits for all writes to be acknowledged before passing acknowledgments on to VMs. This approach incurs a noticeable performance cost, and might be solved by setting the storage to present the SR mount point as an asynchronous mode export. Asynchronous exports acknowledge writes that are not actually on disk. Consider the risks of failure carefully in these situations.

Note:

The NFS server must be configured to export the specified path to all hosts in the pool. If this configuration is not done, the creation of the SR and the plugging of the PBD record fails.

The XenServer NFS implementation uses TCP by default. If your situation allows, you can configure the implementation to use UDP in scenarios where there might be a performance benefit. To do this configuration, when creating an SR, specify the `device-config` parameter `useUDP=true`.

If you are creating an SR for IP-based storage (including NFS), you can configure one of the following as the storage network: the NIC that handles the management traffic or a new NIC for the storage traffic. To assign an IP address to a NIC, see [Configure a dedicated storage NIC](#).

xe CLI

For example, to create a shared NFS SR on `192.168.1.10:/export1`, using any version 4 of NFS that is made available by the filer, use the following command:

```
1  xe sr-create content-type=user \  
2  name-label="shared NFS SR" shared=true \  
3  device-config:server=192.168.1.10 device-config:serverpath=/export1 \  
   type=nfs \  
4  device-config:nfsversion="4"
```

To create a non-shared NFS SR on `192.168.1.10:/export1`, using specifically NFS version 4.0, run the following command:

```
1  xe sr-create host-uuid=host_uuid content-type=user \  
2  name-label="Non-shared NFS SR" \  
3  device-config:server=192.168.1.10 device-config:serverpath=/export1 \  
   type=nfs \  
4  device-config:nfsversion="4.0"
```

In XenCenter®

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.
2. Select **NFS** as the physical storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes a summary of the configuration options that you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description based on SR settings** check box and type in the **Description** box. Click **Next** to continue.
4. On the **Location** page, specify the NFS storage target details:
 - **Share Name** The IP address or DNS name of the server and the path. For example, `server:/path` where `server` is the DNS name or IP address of the server computer, and `path` is the directory used to contain the SR. The NFS server must be configured to export the specified path to all servers in the pool.
 - **Advanced Options** You can enter any additional configuration options here.

- **NFS Version** Select the NFS version used by the SR.

Note:

If the underlying storage array does not support NFSv4, NFSv3 is used to mount the share.

5. Click **Scan** to have the wizard scan for existing NFS SRs in the location you specified.
6. The New Storage Repository wizard lists any existing SRs which are not already attached. You can select an SR from the list and attach it as the new storage repository. click **Reattach an existing SR** and select the SR from the list, then click **Finish**.
7. If no existing SRs are found, simply click **Finish** to complete the new SR configuration and close the wizard.

SMB SR

February 18, 2026

Shares on SMB servers (that support SMB 3) can be used immediately as an SR for virtual disks. VDIs are stored in the Microsoft VHD format only. Also, as these SRs can be shared, VDIs stored on shared SRs allow:

- VMs to be started on any XenServer® hosts in a resource pool
- VMs migrate between XenServer hosts in a resource pool using live migration (without noticeable downtime)

Important:

- Support for SMB3 is limited to the ability to connect to a share using the 3 protocol. Extra features like Transparent Failover depend on feature availability in the upstream Linux kernel are not supported in XenServer 9.
- You can locate your SMB3 filer on a Windows Server system. Do not use a Windows Server system that is located on a VM.
- Clustered SMB is not supported with XenServer.
- SMB storage is available for XenServer Premium Edition customers.
- It is highly recommended that you use a dedicated storage network, using at least two bonded links, ideally to independent network switches with redundant power supplies.
- When using SMB storage, do not remove the share from the storage before detaching the SMB SR.

VDIs stored on file-based SRs are *thinly provisioned*. The image file is allocated as the VM writes data into the disk. This approach has the considerable benefit that the VM image files take up only as much

space on the storage as is required. For example, if a 100 GB VDI is allocated for a VM and an OS is installed, the VDI file only reflects the size of the OS data written to the disk rather than the entire 100 GB.

VHD files may also be chained, allowing two VDIs to share common data. In cases where a file-based VM is cloned, the resulting VMs share the common on-disk data at the time of cloning. Each VM proceeds to make its own changes in an isolated copy-on-write version of the VDI. This feature allows file-based VMs to be quickly cloned from templates, facilitating fast provisioning and deployment of new VMs. The maximum supported length of VHD chains is 30.

File-based SRs and VHD implementations in XenServer assume that they have full control over the SR directory on the file server. Administrators must not modify the contents of the SR directory, as this action can risk corrupting the contents of VDIs.

XenServer has been tuned for enterprise-class storage that uses non-volatile RAM to provide fast acknowledgments of write requests while maintaining a high degree of data protection from failure.

Warning:

As VDIs on file-based SRs are created as thin-provisioned VDIs, administrators must ensure that the file-based SRs have enough disk space for all required VDIs. XenServer hosts do not enforce that the space required for VDIs on file-based SRs is present.

Ensure that you monitor the free space on your SR. If the SR usage grows to 100%, further writes from VMs fail. These failed writes can cause the VM to freeze or crash.

Device-config parameters for SMB SRs:

Parameter Name	Description	Required?
<code>server</code>	Full path to share on server	Yes
<code>username</code>	User account with RW access to share	Optional
<code>password_secret</code>	(Recommended) Secret ID for the password for the user account, which can be used instead of the password.	Optional
<code>password</code>	Password for the user account. We recommend that you use the <code>password_secret</code> parameter instead.	Optional

Create a shared SMB SR (SMB)

To create an SMB SR, provide the host name or IP address of the SMB server, the full path of the exported share, and appropriate credentials.

xe CLI

Note:

When running the `sr-create` command, we recommend that you use the `device-config:password_secret` argument instead of specifying the password on the command line. For more information, see [Secrets](#).

For example, to create a shared SMB SR on `192.168.1.10:/share1`, use the following command:

```
1  xe sr-create content-type=user \
2  name-label="Example shared SMB SR" shared=true \
3  device-config:server=//192.168.1.10/share1 \
4  device-config:username=valid_username device-config:password_secret
   =valid_password_secret type=smb
```

To create a non-shared SMB SR, run the following command:

```
1  xe sr-create host-uuid=host_uuid content-type=user \
2  name-label="Non-shared SMB SR" \
3  device-config:server=//192.168.1.10/share1 \
4  device-config:username=valid_username device-config:password_secret
   =valid_password_secret type=smb
```

In XenCenter®

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar.
2. Select **SMB** as the physical storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description based on SR settings** check box and type in the **Description** box. Click **Next** to continue.
4. On the **Location** page, specify the details of the storage target:
 - **Share Name** The IP address or DNS name of the server and the path. For example, `\\server\path` where `server` is the DNS name or IP address of the server computer, and `path` is a folder or file name. Configure the SMB server to export the specified path to all servers in the pool.

- **User name and Password** (Optional) To connect to an SMB server using a different user name, enter your login user name and password.
5. Click **Scan** to have the wizard scan for existing SMB SRs in the location you specified.
 6. The **New Storage Repository** wizard lists any existing SRs which are not already attached. You can select an SR from the list and attach it as the new storage repository. Click **Reattach an existing SR** and select the SR from the list, then click **Finish**.
 7. If no existing SRs are found, simply click **Finish** to complete the new SR configuration and close the wizard.

Software iSCSI storage

December 17, 2025

XenServer supports shared storage repositories (SRs) on iSCSI LUNs. iSCSI is supported using the Open-iSCSI software iSCSI initiator or by using a supported iSCSI Host Bus Adapter (HBA). The steps for using iSCSI HBAs are identical to the steps for Fibre Channel HBAs. Both sets of steps are described in [Create a Shared LVM over Fibre Channel / Fibre Channel over Ethernet / iSCSI HBA or SAS SR](#).

Shared iSCSI support using the software iSCSI initiator is implemented based on the Linux Volume Manager (LVM). This feature provides the same performance benefits provided by LVM VDIs in the local disk case. Shared iSCSI SRs using the software-based host initiator can support VM agility using live migration: VMs can be started on any XenServer® host in a resource pool and migrated between them with no noticeable downtime.

iSCSI SRs use the entire LUN specified at creation time and cannot span more than one LUN. CHAP support is provided for client authentication, during both the data path initialization and the LUN discovery phases.

Note:

The block size of an iSCSI LUN must be 512 bytes. To use storage with 4 KB physical blocks, the storage must also support emulation of 512-byte allocation blocks (the logical block size must be 512 bytes).

XenServer host iSCSI configuration

All iSCSI initiators and targets must have a unique name to ensure that they can be uniquely identified on the network. An initiator has an iSCSI initiator address, and a target has an iSCSI target address. Collectively these names are called iSCSI Qualified Names, or IQNs.

XenServer hosts support a single iSCSI initiator, which is automatically created and configured with a random IQN during host installation. The single initiator can be used to connect to multiple iSCSI targets concurrently.

iSCSI targets commonly provide access control using iSCSI initiator IQN lists. All iSCSI targets/LUNs that your XenServer host accesses must be configured to allow access by the host's initiator IQN. Similarly, targets/LUNs to be used as shared iSCSI SRs must be configured to allow access by all host IQNs in the resource pool.

Note:

iSCSI targets that do not provide access control typically default to restricting LUN access to a single initiator to ensure data integrity. If an iSCSI LUN is used as a shared SR across multiple hosts in a pool, ensure that multi-initiator access is enabled for the specified LUN.

The XenServer host IQN value can be adjusted using XenCenter®, or using the CLI with the following command when using the iSCSI software initiator:

```
1    xe host-param-set uuid=valid_host_id other-config:iscsi_iqn=
    new_initiator_iqn
```

Warning:

- Each iSCSI target and initiator must have a unique IQN. If a non-unique IQN identifier is used, data corruption or denial of LUN access can occur.
- Do not change the XenServer host IQN with iSCSI SRs attached. Doing so can result in failures connecting to new targets or existing SRs.

If you are creating an SR for IP-based storage (including iSCSI), you can configure one of the following as the storage network: the NIC that handles the management traffic or a new NIC for the storage traffic. To assign an IP address to a NIC, see [Configure a dedicated storage NIC](#).

Shared LVM storage

The Shared LVM type represents disks as Logical Volumes within a Volume Group created on an iSCSI (FC or SAS) LUN.

Note:

The block size of an iSCSI LUN must be 512 bytes. To use storage with 4 KB physical blocks, the storage must also support emulation of 512-byte allocation blocks (the logical block size must be 512 bytes).

Device-config parameters for LVMo iSCSI SRs:

Parameter Name	Description	Required?
target	The IP address or host name of the iSCSI target on the SAN that hosts the SR. This can also be a comma-separated list of values to connect to multiple targets.	Yes
targetIQN	The iSCSI Qualified Name (IQN) of the target on the iSCSI SAN that hosts the SR, or * to connect to all IQNs.	Yes
SCSIid	The SCSI bus ID of the destination LUN	Yes
multihomed	Enable multi-homing to this target	No (defaults to the same value as host.other_config:multipathing)
chapuser	The user name to be used for CHAP authentication	No
chappassword_secret	(Recommended) Secret ID for the password to be used for CHAP authentication. Pass a secret instead of a password.	No
chappassword	The password to be used for CHAP authentication. We recommend that you use the chappassword_secret parameter instead.	No
port	The network port number on which to query the target	No
usediscoverynumber	The specific iSCSI record index to use	No
incoming_chapuser	The user name that the iSCSI filter uses to authenticate against the host	No
incoming_chappassword_secret	(Recommended) Secret ID for the password that the iSCSI filter uses to authenticate against the host.	No

Parameter Name	Description	Required?
<code>incoming_chappassword</code>	The password that the iSCSI filter uses to authenticate against the host. We recommend that you use the <code>incoming_chappassword_secret</code> parameter instead.	No

Create a shared LVM over iSCSI SR by using the Software iSCSI initiator

xe CLI

Note:

When running the `sr-create` command, we recommend that you use the `device-config:chappassword_secret` argument instead of specifying the password on the command line. For more information, see [Secrets](#).

To create a shared LVM over iSCSI SR on a specific LUN of an iSCSI target, use the following command.

```

1  xe sr-create host-uuid=valid_uuid content-type=user \
2  name-label="Example shared LVM over iSCSI SR" shared=true \
3  device-config:target=target_ip= device-config:targetIQN=target_iqn=
   \
4  device-config:SCSIid=scsi_id \
5  type=lvmoiscsi
```

In XenCenter

Note:

Before performing the following steps, ensure the iSCSI Initiator IQN is set appropriately for all hosts in the pool. For more information, see [Changing Server Properties](#).

1. Open the **New Storage Repository** wizard: click **New Storage** on the toolbar. Alternatively:
 - On the **Storage** tab for the selected pool or server, click **New SR**.
 - On the **Storage** menu, click **New SR**.
 - In the **Resources** pane, select a server or pool then right-click and click **New SR** on the shortcut menu.
2. Select **Software iSCSI** as the physical storage type, then click **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes a summary of the configuration options you select as

you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box. Click **Next** to continue.

4. On the **Provisioning** page, select the type of provisioning to use. The options available are
 - Thin provisioning (GFS2). This type of provisioning is only available on GFS2 clustered pools. For more information about GFS2 clustering, see [Change pool properties](#)
 - Full provisioning (LVM)

Click **Next** to continue.

5. On the **Location** page, specify the iSCSI target details:
 - **Target Host:** The IP address or DNS name of the iSCSI target. This can also be a comma-separated list of values.
 - **Use CHAP:** If the iSCSI target is configured to use CHAP authentication, select the **Use CHAP** check box and fill in the details:
 - CHAP User: the CHAP authentication user name credential to apply when connecting to the target.
 - CHAP Secret: the CHAP authentication password credential to apply when connecting to the target.
 - **Target IQN:** To specify the iSCSI target IQN, click the **Discover IQNs** button and then choose an IQN from the **Target IQN** list.

Important:

The iSCSI target and all servers in the pool must not have the same IQN set. Every iSCSI target and initiator must have a unique IQN. If a non-unique IQN identifier is used, data corruption can occur or access to the target can be denied or both.

- **Target LUN:** To specify the LUN on which to create the storage repository, click the **Discover LUNs** button. Choose a LUN from the **Target LUN** list.

Each individual iSCSI storage repository must be contained entirely on a single LUN. The SR cannot span more than one LUN. If the LUN already contains an SR, choose either to use the existing SR or to replace the existing SR with a new one. Replacing the existing SR destroys any data present on the disk.

6. Click **Finish** to complete the new SR configuration and close the wizard.

Hardware HBA LUNs

May 21, 2026

This section covers various operations required to manage SAS, Fibre Channel, and iSCSI HBAs.

LVM over Hardware HBA

The LVM over hardware HBA type represents disks as VHDs on Logical Volumes within a Volume Group created on an HBA LUN that provides, for example, hardware-based iSCSI or FC support.

XenServer® hosts support Fibre Channel SANs through Emulex or QLogic host bus adapters (HBAs). All Fibre Channel configuration required to expose a Fibre Channel LUN to the host must be completed manually. This configuration includes storage devices, network devices, and the HBA within the XenServer host. After all FC configuration is complete, the HBA exposes a SCSI device backed by the FC LUN to the host. The SCSI device can then be used to access the FC LUN as if it were a locally attached SCSI device.

Use the `sr-probe` command to list the LUN-backed SCSI devices present on the host. This command forces a scan for new LUN-backed SCSI devices. The path value returned by `sr-probe` for a LUN-backed SCSI device is consistent across all hosts with access to the LUN. Therefore, this value must be used when creating shared SRs accessible by all hosts in a resource pool.

The same features apply to QLogic iSCSI HBAs.

Note:

XenServer support for Fibre Channel does not support direct mapping of a LUN to a VM. HBA-based LUNs must be mapped to the host and specified for use in an SR. VDIs within the SR are exposed to VMs as standard block devices.

The block size of an LVM over HBA LUN must be 512 bytes. To use storage with 4 KB physical blocks, the storage must also support emulation of 512-byte allocation blocks (the logical block size must be 512 bytes).

Remove HBA-based SAS, FC, or iSCSI device entries

Note:

This step is not required. We recommend that only power users perform this process if it is necessary.

Each HBA-based LUN has a corresponding global device path entry under `/dev/disk/by-scsibus` in the format `<SCSIid>-<adapter>:<bus>:<target>:<lun>` and a standard device path under `/dev`. To remove the device entries for LUNs no longer in use as SRs, use the following steps:

1. Use `sr-forget` or `sr-destroy` as appropriate to remove the SR from the XenServer host database. See [Remove SRs](#) for details.

2. Remove the zoning configuration within the SAN for the desired LUN to the desired host.
3. Use the `sr-probe` command to determine the ADAPTER, BUS, TARGET, and LUN values corresponding to the LUN to be removed. For more information, [Probe an SR](#).
4. Remove the device entries with the following command:

```
1 echo "1" > /sys/class/scsi_device/adapter:bus:target:lun/device/delete
```

Warning:

Make sure that you are certain which LUN you are removing. Accidentally removing a LUN required for host operation, such as the boot or root device, renders the host unusable.

Shared LVM storage

The Shared LVM type represents disks as Logical Volumes within a Volume Group created on an iSCSI (FC or SAS) LUN.

Note:

The block size of an iSCSI LUN must be 512 bytes. To use storage with 4 KB physical blocks, the storage must also support emulation of 512-byte allocation blocks (the logical block size must be 512 bytes).

Device-config parameters for LVMoHBA SRs:

Parameter name	Description	Required?
<code>SCSIid</code>	Device SCSI ID	Yes

Create a Shared LVM over Fibre Channel / Fibre Channel over Ethernet / iSCSI HBA or SAS SR

SRs of type LVMoHBA can be created and managed using the xe CLI or XenCenter.

xe CLI To create a shared LVMoHBA SR, perform the following steps on each host in the pool:

1. Zone in one or more LUNs to each XenServer host in the pool. This process is highly specific to the SAN equipment in use. For more information, see your SAN documentation.
2. If necessary, configure the HBA using third-party tools in the BIOS or firmware.
 - For information about configuring QLogic Fibre Channel and iSCSI HBAs, see the [Marvell](#) website.

- For information about configuring Emulex Fibre Channel HBAs, see the [Broadcom](#) website.
3. Use the `sr-probe` command to determine the global device path of the HBA LUN. The `sr-probe` command forces a rescan of HBAs installed in the system to detect any new LUNs that have been zoned to the host. The command returns a list of properties for each LUN found. Specify the `host-uuid` parameter to ensure that the probe occurs on the desired host.

The global device path returned as the `<path>` property is common across all hosts in the pool. Therefore, this path must be used as the value for the `device-config:device` parameter when creating the SR.

If multiple LUNs are present, use the vendor, LUN size, LUN serial number, or the SCSI ID from the `<path>` property to identify the desired LUN.

```

1  xe sr-probe type=lvmohba \
2  host-uuid=1212c7b3-f333-4a8d-a6fb-80c5b79b5b31
3  Error code: SR_BACKEND_FAILURE_90
4  Error parameters: , The request is missing the device
   parameter, \
5  <?xml version="1.0" ?>
6  <Devlist>
7      <BlockDevice>
8          <path>
9              /dev/disk/by-id/scsi-360
               a9800068666949673446387665336f
10             </path>
11             <vendor>
12                 HITACHI
13             </vendor>
14             <serial>
15                 730157980002
16             </serial>
17             <size>
18                 80530636800
19             </size>
20             <adapter>
21                 4
22             </adapter>
23             <channel>
24                 0
25             </channel>
26             <id>
27                 4
28             </id>
29             <lun>
30                 2
31             </lun>
32             <hba>
33                 qla2xxx
34             </hba>
35         </BlockDevice>
36         <Adapter>

```



```

37         <host>
38             Host4
39         </host>
40         <name>
41             qla2xxx
42         </name>
43         <manufacturer>
44             QLogic HBA Driver
45         </manufacturer>
46         <id>
47             4
48         </id>
49     </Adapter>
50 </Devlist>

```

4. On the pool coordinator, create the SR. Specify the global device path returned in the `<path>` property from `sr-probe`. PBDs are created and plugged for each host in the pool automatically.

```

1  xe sr-create host-uuid=valid_uuid \
2  content-type=user \
3  name-label="Example shared LVM over HBA SR" shared=true \
4  device-config:SCSIid=device_scsi_id type=lvmohba

```

Note:

You can use the XenCenter Repair Storage Repository function to retry the PBD creation and plugging portions of the `sr-create` operation. This function can be valuable in cases where the LUN zoning was incorrect for one or more hosts in a pool when the SR was created. Correct the zoning for the affected hosts and use the Repair Storage Repository function instead of removing and re-creating the SR.

In XenCenter

1. To open the **New Storage Repository** wizard, you can do any of the following actions:
 - On the toolbar, select **New Storage**.
 - On the **Storage** tab for the selected pool or server, select **New SR**.
 - On the **Storage** menu, select **New SR**.
 - In the **Resources** pane, select a server or pool then right-click and select **New SR** on the shortcut menu.
2. Select **Hardware HBA** as the physical storage type and then select **Next**.
3. On the **Name** page, enter the name of the new SR. By default, the wizard generates a description of the SR. This description includes a summary of the configuration options you select as you progress through the wizard. To enter your own description, clear the **Auto-generate description** check box and type in the **Description** box. Click **Next** to continue to the **Provisioning**

page.

4. On the **Provisioning** page, select the provisioning type. The options available are
 - Thin provisioning (GFS2). This type of provisioning is only available on GFS2 clustered pools. For more information about GFS2 clustering, see [Change pool properties](#)
 - Full provisioning (LVM)

Click **Next** to continue to the **Location** page.

The wizard scans for available LUNs and then displays a page listing all the LUNs found. Select a LUN from the list and click **Create**.

Note:

A warning message is displayed if there are existing SRs on the LUN you have selected. Review the details and choose one of the following options.

- To use the existing, click **Reattach**.
- To delete the existing SR and to create an SR, click **Format**.
- If you prefer to select a different LUN, click **Cancel** and select a LUN from the list.

The **Summary** page displays information about the new SR. Read the information and then click **Finish** to complete the SR creation process.

Thin-provisioned shared GFS2 block storage

April 27, 2026

Thin provisioning better utilizes the available storage by allocating disk storage space to VDIs as data is written to the virtual disk, rather than allocating the full virtual size of the VDI in advance. Thin provisioning enables you to significantly reduce the amount of space required on a shared storage array, and with that your Total Cost of Ownership (TCO).

Warning:

GFS2 has a number of limitations and complexities which should be carefully considered before proceeding. In many cases an alternative storage strategy will be appropriate.

Limitations of GFS2 include:

- Not fully featured (no support for DR, no support for storage migration, pool sizes limited to 16)
- Performance limitations, making it unsuitable for large, I/O intensive workloads
- Requires cluster services, which can bring stability issues

- More complex to understand, resulting in increased misconfigurations and errors

Understanding these limitations and complexities, it is only recommended to use GFS2 where no alternatives exist to meet requirements. The scenarios where GFS2 may be appropriate are:

- *Thin provisioning is required and not provided by your SAN.* The recommendation is to use LVM storage with a SAN that provides the thin provisioning requirements. Where this is not possible, GFS2 can be considered.
- *Support for >2TiB disks.* GFS2 will support disks up to 16TB. Where >2TiB disks are required, GFS2 can be considered. Customers should undertake testing to ensure that the environment can meet load and performance requirements.

If your SAN supports thin provisioning, we recommend that you use this, with LVM, in preference to GFS2 unless you need VM disks bigger than 2 TiB. For cases where VM disks larger than 2 TiB are required, consider multiple disks in LVM to achieve this.

The shared GFS2 type represents disks as a filesystem created on an iSCSI or HBA LUN. VDIs stored on a GFS2 SR are stored in the QCOW2 image format.

This article describes how to set up your GFS2 environment by using the xe CLI. To set up a GFS2 environment by using XenCenter, see [the XenCenter product documentation](#).

1. Plan your GFS2 environment

To provide the benefits of thin provisioning on shared block storage without risk of data loss, your pool must deliver a good level of reliability and connectivity. It is crucial that the hosts in the resource pool that uses GFS2 can reliably communicate with one another. To ensure this, XenServer® requires that you use a clustered pool with your GFS2 SR. We also recommend that you design your environment and configure XenServer features to provide as much resiliency and redundancy as possible.

Before setting up your XenServer pool to work with GFS2 SRs, review the following requirements and recommendations for an ideal GFS2 environment:

- **Recommended:** Configure redundant networking infrastructure.
- **Recommended:** Create a dedicated bonded network
- **Required:** Set up a clustered pool
- **Optional** Increase your control domain memory
- **Recommended:** Configure storage multipathing
- **Required:** Create a GFS2 SR

A clustered pool with GFS2 SRs has some differences in behavior to other types of pool and SR. For more information, see Constraints.

2. Configure redundant networking infrastructure

A bonded network links two or more NICs together to create a single channel for network traffic. We recommend that you use a bonded network for your clustered pool traffic. However, before you set up your bonded network, ensure that your network hardware configuration promotes redundancy in the bonded network. Consider implementing as many of these recommendations as is feasible for your organization and environment.

The following best practices add resiliency against software, hardware, or power failures that can affect your network switches.

- Ensure that you have separate physical network switches available for use in the bonded network, not just ports on the same switch.
- Ensure that the separate switches draw power from different, independent power distribution units (PDUs).
- If possible, in your data center, place the PDUs on different phases of the power feed or even feeds provided by different utility companies.
- Consider using uninterruptible power supply units to ensure that the network switches and servers can continue to function or perform an orderly shutdown in the event of a power failure.

3. Create a dedicated bonded network

It is important to ensure that hosts in a clustered pool can communicate reliably with one another. Creating a bonded network for this pool traffic increases the resiliency of your clustered pool.

Note:

The cluster network cannot be on a non-management VLAN.

A bonded network creates a bond between two or more NICs to create a single, high-performing channel that your clustered pool can use for cluster heartbeat traffic. We strongly recommend that this bonded network is not used for any other traffic. Create a separate network for the pool to use for management traffic.

Warning:

If you choose not to follow this recommendation, you are at a higher risk of losing cluster management network packets. Loss of cluster management network packets can cause your clustered pool to lose quorum and some or all hosts in the pool will self-fence.

If your cluster is fencing or facing a problem in this unrecommended configuration, XenServer Support might ask you to reproduce the same problem on a recommended configuration during the course of investigation.

To create a bonded network to use as the GFS2 cluster network:

1. If you have a firewall between the hosts in your pool, ensure that hosts can communicate on the cluster network using the following ports:

- TCP: 8892, 8896, 21064
- UDP: 5404, 5405

For more information, see [Communication ports used by XenServer](#).

2. Open a console on the XenServer host that you want to act as the pool coordinator.
3. Create a network for use with the bonded NIC by using the following command:

```
1 xe network-create name-label=bond0
```

The UUID of the new network is returned.

4. Find the UUIDs of the PIFs to use in the bond by using the following command:

```
1 xe pif-list
```

5. Create your bonded network in either active-active mode, active-passive mode, or LACP bond mode. Depending on the bond mode you want to use, complete one of the following actions:

- To configure the bond in active-active mode (default), use the `bond-create` command to create the bond. Using commas to separate the parameters, specify the newly created network UUID and the UUIDs of the PIFs to be bonded:

```
1 xe bond-create network-uuid=<network_uuid> /  
2   pif-uuids=<pif_uuid_1>,<pif_uuid_2>,<pif_uuid_3>,<  
   pif_uuid_4>
```

Type two UUIDs when you are bonding two NICs and four UUIDs when you are bonding four NICs. The UUID for the bond is returned after running the command.

- To configure the bond in active-passive or LACP bond mode, use the same syntax, add the optional `mode` parameter, and specify `lACP` or `active-backup`:

```
1 xe bond-create network-uuid=<network_uuid> /  
2   pif-uuids=<pif_uuid_1>,<pif_uuid_2>,<pif_uuid_3>,<  
   pif_uuid_4> /  
3   mode=balance-slb | active-backup | lACP
```

After you create your bonded network on the pool coordinator, when you join other XenServer hosts to the pool, the network and bond information is automatically replicated to the joining server.

For more information, see [Networking](#).

Note:

- Changing the IP address of the cluster network by using XenCenter® requires clustering and GFS2 to be temporarily disabled.
- Do not change the bonding of your cluster network while the cluster is live and has running VMs. This action can cause hosts in the cluster to hard restart (fence).
- If you have an IP address conflict (multiple hosts having the same IP address) on your cluster network involving at least one host with clustering enabled, the cluster does not form correctly and the hosts are unable to fence when required. To fix this issue, resolve the IP address conflict.

The cluster timeout value of your pool depends on how many hosts are in your cluster. Run the following command to find the `token-timeout` value in seconds for the pool:

```
1 xe cluster-param-get uuid=<cluster_uuid> param-name=token-timeout
```

If the network bond failover time is likely to be greater than the timeout value, your network infrastructure and configuration might not be reliable enough to support a clustered pool.

4. Set up a clustered pool

To use shared GFS2 storage, the XenServer resource pool must be a clustered pool. Enable clustering on your pool before creating a GFS2 SR.

A clustered pool is a pool of XenServer hosts that are more closely connected and coordinated than hosts in non-clustered pools. The hosts in the cluster maintain constant communication with each other on a selected network. All hosts in the cluster are aware of the state of every host in the cluster. This host coordination enables the cluster to control access to the contents of the GFS2 SR. To ensure that the clustered pool always remains in communication, each host in a cluster must always be in communication with at least half of the hosts in the cluster (including itself). This state is known as a host having quorum. If a host does not have quorum, it hard restarts and removes itself from the cluster. This action is referred to as ‘fencing’.

For more information, see [Clustered pools](#).

Before you start setting up your clustered pool, ensure that the following prerequisites are met:

- Plan to create a pool of between 3 and 16 hosts.

Where possible, use an odd number of hosts in a clustered pool as this ensures that hosts are always able to determine if they have quorum. We recommend that you use clustering only in pools containing at least three hosts, as pools of two hosts are sensitive to self-fencing the entire pool.

Clustered pools only support up to 16 hosts per pool.

- All XenServer hosts in the clustered pool must have at least 2 GiB of control domain memory.
- All hosts in the cluster must use static IP addresses for the cluster network.
- If you are clustering an existing pool, ensure that high availability is disabled. You can enable high availability again after clustering is enabled.

To use the xe CLI to create a clustered pool:

1. Create a resource pool of at least three XenServer hosts.

Repeat the following steps on each joining XenServer host that is not the pool coordinator:

- a) Open a console on the XenServer host.
- b) Join the XenServer host to the pool on the pool coordinator by using the following command:

```
1 xe pool-join master-address=<master_address> /
2   master-username=<administrators_username> /
3   master-password=<password>
```

The value of the `master-address` parameter must be set to the fully qualified domain name of the XenServer host that is the pool coordinator. The `password` must be the administrator password set when the pool coordinator was installed.

For more information, see [Hosts and resource pools](#).

2. For every PIF that belongs to this network, set `disallow-unplug=true`.

- a) Find the UUIDs of the PIFs that belong to the network by using the following command:

```
1 xe pif-list
```

- b) Run the following command on a XenServer host in your resource pool:

```
1 xe pif-param-set disallow-unplug=true uuid=<pif_uuid>
```

3. Enable clustering on your pool. Run the following command on a XenServer host in your resource pool:

```
1 xe cluster-pool-create network-uuid=<network_uuid>
```

Provide the UUID of the bonded network that you created in an earlier step.

5. Increase your control domain memory

If you have insufficient control domain memory on your hosts, your pool can experience network instability. Network instability can cause problems for a clustered pool with GFS2 SRs.

It is important to ensure that your clustered pool has an appropriate amount of control domain memory. For information about changing the amount of control domain memory and monitoring the memory behavior, see [Memory usage](#).

6. Configure storage multipathing

Ensure that storage multipathing is set up between your clustered pool and your GFS2 SR.

Multipathing routes storage traffic to a storage device over multiple paths for redundancy. All routes can have active traffic on them during normal operation, which results in increased throughput.

Before enabling multipathing, verify that the following statements are true:

- Your ethernet or fibre switch is configured to make multiple targets available on your storage server.

For example, an iSCSI storage back-end queried for `sendtargets` on a given portal returns multiple targets, as in the following example:

```
1  iscsiadm -m discovery --type sendtargets --portal 192.168.0.161
2  192.168.0.161:3260,1 iqn.strawberry:litchie
3  192.168.0.204:3260,2 iqn.strawberry:litchie
```

However, you can perform additional configuration to enable iSCSI multipath for arrays that only expose a single target. For more information, see [iSCSI multipath for arrays that only expose a single target](#).

- For iSCSI only, the control domain (dom0) has an IP address on each subnet used by the multipathed storage.

Ensure that for each path to the storage, you have a NIC and that there is an IP address configured on each NIC. For example, if you want four paths to your storage, you must have four NICs that each have an IP address configured.

- For iSCSI only, every iSCSI target and initiator has a unique IQN.
- For iSCSI only, the iSCSI target ports are operating in portal mode.
- For HBA only, multiple HBAs are connected to the switch fabric.
- If possible, use multiple redundant switches.

To enable multipathing by using the xe CLI

We recommend that you enable multipathing for all hosts in your pool *before* creating the SR. If you create the SR before enabling multipathing, you must put your hosts into maintenance mode to enable multipathing.

1. Open a console on the XenServer host.
2. Unplug all PBDs on the host by using the following command:

```
1 xe pbd-unplug uuid=<pbd_uuid>
```

You can use the command `xe pbd-list` to find the UUID of the PBDs.

3. Set the value of the `multipathing` parameter to `true` by using the following command:

```
1 xe host-param-set uuid=<host uuid> multipathing=true
```

4. If there are existing SRs on the hosts running in single path mode that have multiple paths:
 - Migrate or suspend any running guests with virtual disks in the affected SRs.
 - Replug the PBD of any affected SRs to reconnect them using multipathing:

```
1 xe pbd-plug uuid=<pbd_uuid>
```

5. Repeat these steps to enable multipathing on all hosts in the pool.

Ensure that you enable multipathing on all hosts in the pool. All cabling and, in the case of iSCSI, subnet configurations must match the corresponding NICs on each host.

For more information, see [Storage multipathing](#).

7. Create a GFS2 SR

Create your shared GFS2 SR on an iSCSI or an HBA LUN that is visible to all XenServer hosts in your resource pool. We do not recommend using a thin-provisioned LUN with GFS2. However, if you do choose this configuration, you must ensure that the LUN always has enough space to allow XenServer to write to it.

You can add up to 62 GFS2 SRs to a clustered pool.

If you have previously used your block-based storage device for thick provisioning with LVM, this is detected by XenServer. XenCenter gives you the opportunity to use the existing LVM partition or to format the disk and set up a GFS2 partition.

Create a shared GFS2 over iSCSI SR

You can use the `xe` CLI to create a GFS2 over iSCSI SR.

Device-config parameters for GFS2 SRs:

Parameter Name	Description	Required?
<code>provider</code>	The block provider implementation. In this case, <code>iscsi</code> .	Yes
<code>target</code>	The IP address or hostname of the iSCSI filer that hosts	Yes
<code>targetIQN</code>	The IQN target of iSCSI filer that hosts the SR	Yes
<code>SCSIid</code>	Device SCSI ID	Yes

You can find the values to use for these parameters by using the `xe sr-probe-ext` command.

```
1 xe sr-probe-ext type=<type> host-uuid=<host_uuid> device-config:=<config> sm-config:=<sm_config>
```

1. Start by running the following command:

```
1 xe sr-probe-ext type=gfs2 device-config:provider=iscsi
```

The output from the command prompts you to supply additional parameters and gives a list of possible values at each step.

2. Repeat the command, adding new parameters each time.
3. When the command output starts with `Found the following complete configurations that can be used to create SRs:`, you can locate the SR by using the `xe sr-create` command and the `device-config` parameters that you specified.

Example output:

```
1 Found the following complete configurations that can be used to
  create SRs:
2 Configuration 0:
3   SCSIid      : 36001405852f77532a064687aea8a5b3f
4   targetIQN: iqn.2009-01.example.com:iscsi192a25d6
5   target: 198.51.100.27
6   provider: iscsi
7
8
9 Configuration 0 extra information:
```

To create a shared GFS2 SR on a specific LUN of an iSCSI target, run the following command on a server in your clustered pool:

```
1 xe sr-create type=gfs2 name-label="Example GFS2 SR" --shared \
2   device-config:provider=iscsi device-config:targetIQN=<target_iqns> \
```

```
3 device-config:target=<portal_address> device-config:SCSIid=<scsci_id>
```

If the iSCSI target is not reachable while GFS2 filesystems are mounted, some hosts in the clustered pool might hard restart (fence).

For more information about working with iSCSI SRs, see [Software iSCSI storage](#).

Create a shared GFS2 over HBA SR

You can use the xe CLI to create a GFS2 over HBA SR.

Device-config parameters for GFS2 SRs:

Parameter name	Description	Required?
<code>provider</code>	The block provider implementation. In this case, <code>hba</code> .	Yes
<code>SCSIid</code>	Device SCSI ID	Yes

You can find the values to use for the SCSIid parameter by using the `xe sr-probe-ext` command.

```
1 xe sr-probe-ext type=<type> host-uuid=<host_uuid> device-config:=<config> sm-config:=<sm_config>
```

1. Start by running the following command:

```
1 xe sr-probe-ext type=gfs2 device-config:provider=hba
```

The output from the command prompts you to supply additional parameters and gives a list of possible values at each step.

2. Repeat the command, adding new parameters each time.
3. When the command output starts with `Found the following complete configurations that can be used to create SRs:`, you can locate the SR by using the `xe sr-create` command and the `device-config` parameters that you specified.

Example output:

```
1 Found the following complete configurations that can be used to
  create SRs:
2 Configuration 0:
3   SCSIid       : 36001405852f77532a064687aea8a5b3f
4   targetIQN: iqn.2009-01.example.com:iscsi192a25d6
```

```
5         target: 198.51.100.27
6         provider: iscsi
7
8
9 Configuration 0 extra information:
```

To create a shared GFS2 SR on a specific LUN of an HBA target, run the following command on a server in your clustered pool:

```
1 xe sr-create type=gfs2 name-label="Example GFS2 SR" --shared \
2   device-config:provider=hba device-config:SCSIid=<device_scsi_id>
```

For more information about working with HBA SRs, see [Hardware HBA storage](#).

What's next?

Now that you have your GFS2 environment set up, it is important that you maintain the stability of your clustered pool by ensuring it has quorum. For more information, see [Manage your clustered pool](#).

If you encounter issues with your GFS2 environment, see [Troubleshoot clustered pools](#).

You can manage your GFS2 SR the same way as you do other SRs. For example, you can add capacity to the storage array to increase the size of the LUN. For more information, see [Live LUN expansion](#).

Constraints

Shared GFS2 storage currently has the following constraints:

- Intellicache™ is not supported for VMs using a GFS2 SR.
- As with any thin-provisioned SR, if the GFS2 SR usage grows to 100%, further writes from VMs fail. These failed writes can then lead to failures within the VM, possible data corruption, or both.
- XenCenter shows an alert when your SR usage grows to 80%. Ensure that you monitor your GFS2 SR for this alert and take the appropriate action if seen. On a GFS2 SR, high usage causes a performance degradation. We recommend that you keep your SR usage below 80%.
- VM migration with storage migration (live or offline) is not supported for VMs whose VDIs are on a GFS2 SR. You also cannot migrate VDIs from another type of SR to a GFS2 SR.
- The Software FCoE transport is not supported with GFS2 SRs (for fully offloaded FCoE use HBA).
- Trim/unmap is not supported on GFS2 SRs.
- CHAP is not supported on GFS2 SRs.

- You cannot export VDIs that are greater than 2 TiB as VHD or OVA/OVF. However, you can export VMs with VDIs larger than 2 TiB in XVA format.
- We do not recommend using a thin-provisioned LUN with GFS2. However, if you do choose this configuration, you must ensure that the LUN always has enough space to allow XenServer to write to it.
- We do not recommend using SAN deduplication with GFS2 SRs. However, if you do choose this configuration, you must use suitable external monitoring of your SAN utilization to ensure that there is always space for XenServer to write to.
- Your GFS2 file system cannot be larger than 100 TiB.
- You cannot have more than 62 GFS2 SRs in your pool.
- Clustered pools only support up to 16 hosts per pool.
- To enable HA on your clustered pool, the heartbeat SR must be a GFS2 SR.
- For cluster traffic, we strongly recommend that you use a bonded network that uses at least two different network switches. Do not use this network for any other purposes.
- Changing the IP address of the cluster network by using XenCenter requires clustering and GFS2 to be temporarily disabled.
- Do not change the bonding of your cluster network while the cluster is live and has running VMs. This action can cause hosts in the cluster to hard restart (fence).
- If you have an IP address conflict (multiple hosts having the same IP address) on your cluster network involving at least one host with clustering enabled, the cluster does not form correctly and the hosts are unable to fence when required. To fix this issue, resolve the IP address conflict.

Manage storage repositories

June 5, 2026

This section covers various operations required in the ongoing management of Storage Repositories (SRs), including Live VDI Migration.

Remove SRs

A Storage Repository (SR) can be removed either temporarily or permanently.

Detach: Breaks the association between the storage device and the pool or host (PBD Unplug). The SR (and its VDIs) becomes inaccessible. The contents of the VDIs and the meta-information used by

VMs to access the VDIs are preserved. Detach can be used when you temporarily take an SR offline, for example, for maintenance. A detached SR can later be reattached.

Forget: Preserves the contents of the SR on the physical disk, but the information that connects a VM to its VDIs is permanently deleted. For example, allows you to reattach the SR, to another XenServer® host, without removing any of the SR contents.

Destroy: Deletes the contents of the SR from the physical disk.

Note:

When using SMB storage, do not remove the share from the storage before detaching the SMB SR.

For Destroy or Forget, the PBD connected to the SR must be unplugged from the host.

1. Unplug the PBD to detach the SR from the corresponding XenServer host:

```
1 xe pbd-unplug uuid=pbd_uuid
```

2. Use the `sr-destroy` command to remove an SR. The command destroys the SR, deletes the SR and corresponding PBD from the XenServer host database and deletes the SR contents from the physical disk:

```
1 xe sr-destroy uuid=sr_uuid
```

3. Use the `sr-forget` command to forget an SR. The command removes the SR and corresponding PBD from the XenServer host database but leaves the actual SR content intact on the physical media:

```
1 xe sr-forget uuid=sr_uuid
```

Note:

It can take some time for the software object corresponding to the SR to be garbage collected.

Introduce an SR

To reintroduce a previously *forgotten* SR, create a PBD. Manually plug the PBD to the appropriate XenServer hosts to activate the SR.

The following example introduces an SR of type `lvmoiscsi`.

1. Probe the existing SR to determine its UUID:

```
1 xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10 \  
2     device-config:targetIQN=192.168.1.10:filer1 \  
3     device-config:SCSIid=1494554000000000000000000020000000  
    b70200000f000000
```

2. Introduce the existing SR UUID returned from the `sr-probe` command. The UUID of the new SR is returned:

```
1 xe sr-introduce content-type=user name-label="Example Shared LVM  
  over iSCSI SR" \  
2   shared=true uuid=valid_sr_uuid type=lvmoiscsi
```

3. Create a PBD to accompany the SR. The UUID of the new PBD is returned:

```
1 xe pbd-create type=lvmoiscsi host-uuid=valid_uuid sr-uuid=  
  valid_sr_uuid \  
2   device-config:target=192.168.0.1 \  
3   device-config:targetIQN=192.168.1.10:filer1 \  
4   device-config:SCSIId=149455400000000000000000000000002000000  
    b70200000f000000
```

4. Plug the PBD to attach the SR:

```
1 xe pbd-plug uuid=pbd_uuid
```

5. Verify the status of the PBD plug. If successful, the `currently-attached` property is true:

```
1 xe pbd-list sr-uuid=sr_uuid
```

Note:

Perform steps 3 through 5 for each host in the resource pool. These steps can also be performed using the Repair Storage Repository function in XenCenter.

Live LUN expansion

To fulfill capacity requirements, you may need to add capacity to the storage array to increase the size of the LUN provisioned to the XenServer host. Live LUN Expansion allows you to increase the size of the LUN without any VM downtime.

After adding more capacity to your storage array, enter,

```
1 xe sr-scan sr-uuid=sr_uuid
```

This command rescans the SR, and any extra capacity is added and made available.

This operation is also available in XenCenter. Select the SR to resize, and then click **Rescan**.

Warnings:

- It is not possible to shrink or truncate LUNs. Reducing the LUN size on the storage array can lead to data loss.
- Resizing a GFS2 SR can take a long time. While the resize is happening, you might not be

able to add more hosts to the pool or perform other SR operations.

Live VDI migration

Live VDI migration allows the administrator to relocate the VMs Virtual Disk Image (VDI) without shutting down the VM. This feature enables administrative operations such as:

- Moving a VM from cheap local storage to fast, resilient, array-backed storage.
- Moving a VM from a development to production environment.
- Moving between tiers of storage when a VM is limited by storage capacity.
- Performing storage array upgrades.

Limitations and caveats

Live VDI Migration is subject to the following limitations and caveats

- There must be sufficient disk space available on the target repository.

To move virtual disks by using XenCenter

1. In the **Resources** pane, select the SR where the Virtual Disk is stored and then click the **Storage** tab.
2. In the **Virtual Disks** list, select the Virtual Disk that you would like to move, and then click **Move**.
3. In the **Move Virtual Disk** dialog box, select the target SR that you would like to move the VDI to.

Note:

Ensure that the SR has sufficient space for another virtual disk: the available space is shown in the list of available SRs.

4. Click **Move** to move the virtual disk.

For xe CLI reference, see [vdi-pool-migrate](#).

Cold VDI migration between SRs (offline migration)

VDIs associated with a VM can be copied from one SR to another to accommodate maintenance requirements or tiered storage configurations. XenCenter enables you to copy a VM and all of its VDIs to the same or a different SR. A combination of XenCenter and the xe CLI can be used to copy individual VDIs.

For xe CLI reference, see [vm-migrate](#).

Copy all of a VM's VDIs to a different SR

The XenCenter Copy VM function creates copies of all VDIs for a selected VM on the same or a different SR. The source VM and VDIs are not affected by default. To move the VM to the selected SR rather than creating a copy, select the Remove original VM option in the Copy Virtual Machine dialog box.

1. Shut down the VM.
2. Within XenCenter, select the VM and then select the **VM > Copy VM** option.
3. Select the desired target SR.

Copy individual VDIs to a different SR

A combination of the xe CLI and XenCenter can be used to copy individual VDIs between SRs.

1. Shut down the VM.
2. Use the xe CLI to identify the UUIDs of the VDIs to be moved. If the VM has a DVD drive, its `vdi -uuid` is listed as `not in database` and can be ignored.

```
1 xe vbd-list vm-uuid=valid_vm_uuid
```

Note:

The `vbd-list` command displays both the VBD and VDI UUIDs. Be sure to record the VDI UUIDs rather than the VBD UUIDs.

3. In XenCenter, select the **VM Storage** tab. For each VDI to be moved, select the VDI and click the **Detach** button. This step can also be done using the `vbd-destroy` command.

Note:

If you use the `vbd-destroy` command to detach the VDI UUIDs, first check if the VBD has the parameter `other-config:owner` set to `true`. Set this parameter to `false`. Issuing the `vbd-destroy` command with `other-config:owner=true` also destroys the associated VDI.

4. Use the `vdi-copy` command to copy each of the VM VDIs to be moved to the desired SR.

```
1 xe vdi-copy uuid=valid_vdi_uuid sr-uuid=valid_sr_uuid
```

5. In XenCenter, select the **VM Storage** tab. Click the **Attach** button and select the VDIs from the new SR. This step can also be done use the `vbd-create` command.
6. To delete the original VDIs, select the **Storage** tab of the original SR in XenCenter. The original VDIs are listed with an empty value for the VM field. Use the **Delete** button to delete the VDI.

Convert local Fibre Channel SRs to shared SRs

Use the xe CLI and the XenCenter **Repair Storage Repository** feature to convert a local FC SR to a shared FC SR:

1. Upgrade all hosts in the resource pool to XenServer 9.
2. Ensure that all hosts in the pool have the SR's LUN zoned appropriately. See [Probe an SR](#) for details on using the `sr-probe` command to verify that the LUN is present on each host.
3. Convert the SR to shared:

```
1 xe sr-param-set shared=true uuid=local_fc_sr
```

4. The SR is moved from the host level to the pool level in XenCenter, indicating that it is now shared. The SR is marked with a red exclamation mark to show that it is not currently plugged on all hosts in the pool.
5. Select the SR and then select the **Storage > Repair Storage Repository** option.
6. Click **Repair** to create and plug a PBD for each host in the pool.

Reclaim space for block-based storage on the backing array using discard

You can use space reclamation to free up unused blocks on a thinly provisioned LUN. After the space is released, the storage array can then reuse this reclaimed space.

Note:

Space reclamation is only available on some types of storage arrays. To determine whether your array supports this feature and whether it needs a specific configuration, see the [Hardware Compatibility List](#) and your storage vendor specific documentation.

To reclaim the space by using XenCenter:

1. Select the **Infrastructure** view, and then choose the host or pool connected to the SR.
2. Click the **Storage** tab.
3. Select the SR from the list, and click **Reclaim freed space**.
4. Click **Yes** to confirm the operation.
5. Click **Notifications** and then **Events** to view the status of the operation.

For more information, press **F1** in XenCenter to access the Online Help.

To reclaim space by using the xe CLI, you can use the following command:

```
1 xe host-call-plugin host-uuid=host_uuid \  
2   plugin=trim fn=do_trim args:sr_uuid=sr_uuid
```

Notes:

- The operation is only available for LVM-based SRs that are based on thinly provisioned LUNs on the array. Local SSDs can also benefit from space reclamation.
- Space reclamation is not required for file-based SRs such as NFS and EXT3/EXT4. The **Re-claim Freed Space** button is not available in XenCenter for these SR types.
- If you run the space reclamation `xe` command for a file-based SR or a thick-provisioned LVM-based SR, the command returns an error.
- Space reclamation is an intensive operation and can lead to a degradation in storage array performance. Therefore, only initiate this operation when space reclamation is required on the array. We recommend that you schedule this work outside of peak array demand hours.

Automatically reclaim space when deleting snapshots

When deleting snapshots with XenServer, space allocated on LVM-based SRs is reclaimed automatically and a VM reboot is not required. This operation is known as ‘online coalescing’. Online coalescing applies to all types of SR.

In certain cases, automated space reclamation might be unable to proceed. We recommend that you use the offline coalesce tool in these scenarios:

- Under conditions where a VM I/O throughput is considerable
- In conditions where space is not being reclaimed after a period

Notes:

- Running the offline coalesce tool incurs some downtime for the VM, due to the suspend/resume operations performed.
- Before running the tool, delete any snapshots and clones you no longer want. The tool reclaims as much space as possible given the remaining snapshots/clones. If you want to reclaim the entire space, delete all snapshots and clones.
- VM disks must be either on shared or local storage for a single host. VMs with disks in both types of storage cannot be coalesced.

Reclaim space by using the offline coalesce tool

Enable the hidden objects using XenCenter. Click **View > Hidden** objects. In the Resource pane, select the VM for which you want to obtain the UUID. The UUID is displayed in the **General** tab.

In the **Resource** pane, select the resource pool coordinator (the first host in the list). The **General** tab displays the UUID. If you are not using a resource pool, select the VM's host.

1. Open a console on the host and run the following command:

```
1 xe host-call-plugin host-uuid=host-UUID \  
2   plugin=coalesce-leaf fn=leaf-coalesce args:vm-uuid=VM-UUID
```

For example, if the VM UUID is `9bad4022-2c2d-dee6-abf5-1b6195b1dad5` and the host UUID is `b8722062-de95-4d95-9baa-a5fe343898ea`, run the following command:

```
1 xe host-call-plugin host-uuid=b8722062-de95-4d95-9baa-a5fe343898ea \  
2   \plugin=coalesce-leaf fn=leaf-coalesce args:vm-uuid=9bad4022-2  
   c2d-dee6-abf5-1b6195b1dad5
```

2. This command suspends the VM (unless it is already powered down), initiates the space reclamation process, and then resumes the VM.

Notes:

We recommend that you shut down or suspend the VM manually before running the off-line coalesce tool. You can shut down or suspend the VM using either XenCenter or the XenServer CLI. If you run the coalesce tool on a running VM, the tool automatically suspends the VM, performs the required VDI coalesce operations, and resumes the VM. Agile VMs might restart on a different host.

If the Virtual Disk Images (VDIs) to be coalesced are on shared storage, you must run the off-line coalesce tool on the pool coordinator.

If the VDIs to be coalesced are on local storage, run the off-line coalesce tool on the host to which the local storage is attached.

Working with disk I/O

You can configure the disk I/O scheduler and the disk I/O priority settings to change the performance of your disks.

Note:

The disk I/O capabilities described in this section do not apply to GFS2, NFS, or SMB3 storage.

Adjust the disk I/O scheduler

For general performance, the default disk scheduler `none` is applied on all new SR types. The `none` scheduler provides the fairest performance for competing VMs accessing the same device.

1. Adjust the disk scheduler by using the following command:

```
1 xe sr-param-set other-config:scheduler=<option> uuid=<sr_uuid>
```

The value of `<option>` can be one of the following terms: `none`, `bfq`, or `mq-deadline`.

2. Unplug and replug the corresponding PBD for the scheduler parameter to take effect.

```
1 xe pbd-unplug uuid=<pbd_uuid>
2 xe pbd-plug uuid=<pbd_uuid>
```

To apply disk I/O request prioritization, override the default setting and assign the `bfq` disk scheduler to the SR.

Virtual disk I/O request prioritization

Virtual disks have optional I/O request priority settings. You can use these settings to prioritize I/O to a particular VM's disk over others.

Before configuring any disk I/O request priority parameters for a VBD, ensure that the disk scheduler for the SR has been set appropriately. The scheduler parameter must be set to `bfq` on the SR and the associated PBD unplugged and replugged. For information about how to adjust the scheduler, see [Adjusting the disk I/O scheduler](#).

For shared SR, where multiple hosts are accessing the same LUN, the priority setting is applied to VBDs accessing the LUN from the same host. These settings are not applied across hosts in the pool.

The host issues a request to the remote storage, but the request prioritization is done by the remote storage.

Setting disk I/O request parameters These settings can be applied to existing virtual disks by using the `xe vbd-param-set` command with the following parameters:

- `qos_algorithm_type` - This parameter must be set to the value `ionice`, which is the only algorithm supported for virtual disks.
- `qos_algorithm_param` - Use this parameter to set key-value pairs. For virtual disks, `qos_algorithm_param` takes a `sched` key, and depending on the value, also requires a `class` key.

The key `qos_algorithm_param:sched` can have one of the following values:

- `sched=rt` or `sched=real-time` - This value sets the scheduling parameter to real time priority, which requires a `class` parameter to set a value.
- `sched=idle` - This value sets the scheduling parameter to idle priority, which requires no `class` parameter to set any value.

- `sched=anything` - This value sets the scheduling parameter to best-effort priority, which requires a `class` parameter to set a value.

The key `qos_algorithm_param: class` can have one of the following values:

- One of the following keywords: `highest`, `high`, `normal`, `low`, `lowest`.
- An integer between 0 and 7, where 7 is the highest priority and 0 is the lowest. For example, I/O requests with a priority of 5, are given priority over I/O requests with a priority of 2.

Example For example, the following CLI commands set the virtual disk's VBD to use real time priority 5:

```
1 xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_type=ionice
2 xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_params:sched=rt
3 xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_params:class=5
4 xe sr-param-set uuid=<sr_uuid> other-config:scheduler=bfq
5 xe pbd-unplug uuid=<pbd_uuid>
6 xe pbd-plug uuid=<pbd_uuid>
```

Storage multipathing

December 17, 2025

Dynamic multipathing support is available for Fibre Channel and iSCSI storage back-ends.

XenServer® uses Linux native multipathing (DM-MP), the generic Linux multipathing solution, as its multipath handler. However, XenServer supplements this handler with additional features so that XenServer can recognize vendor-specific features of storage devices.

Configuring multipathing provides redundancy for remote storage traffic if there is partial connectivity loss. Multipathing routes storage traffic to a storage device over multiple paths for redundancy and increased throughput. You can use up to 16 paths to a single LUN. Multipathing is an active-active configuration. By default, multipathing uses either round-robin or multibus load balancing depending on the storage array type. All routes have active traffic on them during normal operation, which results in increased throughput.

Important:

We recommend that you enable multipathing for all hosts in your pool *before* creating the SR. If you create the SR before enabling multipathing, you must put your hosts into maintenance mode to enable multipathing.

NIC bonding can also provide redundancy for storage traffic. For iSCSI storage, we recommend configuring multipathing instead of NIC bonding whenever possible.

Multipathing is not effective in the following scenarios:

- NFS storage devices
- You have limited number of NICs and need to route iSCSI traffic and file traffic (NFS or SMB) over the same NIC

In these cases, consider using NIC bonding instead. For more information about NIC bonding, see [Networking](#).

Prerequisites

Before enabling multipathing, verify that the following statements are true:

- Multiple targets are available on your storage server.

For example, an iSCSI storage back-end queried for `sendtargets` on a given portal returns multiple targets, as in the following example:


```
1 iscsiadm -m discovery --type sendtargets --portal 192.168.0.161
2 192.168.0.161:3260,1 iqn.strawberry:litchie
3 192.168.0.204:3260,2 iqn.strawberry:litchie
```

However, you can perform additional configuration to enable iSCSI multipath for arrays that only expose a single target. For more information, see [iSCSI multipath for arrays that only expose a single target](#).

- For iSCSI only, the control domain (dom0) has an IP address on each subnet used by the multipathed storage.

Ensure that for each path you want to have to the storage, you have a NIC and that there is an IP address configured on each NIC. For example, if you want four paths to your storage, you must have four NICs that each have an IP address configured.

- For iSCSI only, every iSCSI target and initiator has a unique IQN.
- For iSCSI only, the iSCSI target ports are operating in portal mode.
- For HBA only, multiple HBAs are connected to the switch fabric.
- When you are configuring secondary interfaces, each secondary interface must be on a separate subnet. For example, if you want to configure two more secondary interfaces for storage, you require IP addresses on three different subnets—one subnet for the management interface, one subnet for Secondary Interface 1, and one subnet for Secondary Interface 2.



This diagram shows how both NICs on the host in a multipath iSCSI configuration must be on different subnets. In this diagram, NIC 1 on the host along with Switch 1 and NIC 1 on both storage controllers are on a different subnet than NIC2, Switch 2, and NIC 2 on the storage controllers.

Enable multipathing

You can enable multipathing in XenCenter® or on the xe CLI.

To enable multipathing by using XenCenter

1. In the XenCenter **Resources** pane, right-click on the host and choose **Enter Maintenance Mode**.
2. Wait until the host reappears in the **Resources** pane with the maintenance mode icon (a blue square) before continuing.
3. On the **General** tab for the host, click **Properties** and then go to the **Multipathing** tab.
4. To enable multipathing, select the **Enable multipathing on this server** check box.
5. Click **OK** to apply the new setting. There is a short delay while XenCenter saves the new storage configuration.
6. In the **Resources** pane, right-click on the host and choose **Exit Maintenance Mode**.
7. Repeat these steps to enable multipathing on all hosts in the pool.

Ensure that you enable multipathing on all hosts in the pool. All cabling and, in the case of iSCSI, subnet configurations must match the corresponding NICs on each host.

To enable multipathing by using the xe CLI

1. Open a console on the XenServer host.
2. Unplug all PBDs on the host by using the following command:

```
1 xe pbd-unplug uuid=<pbd_uuid>
```

You can use the command `xe pbd-list` to find the UUID of the PBDs.

3. Set the value of the `multipathing` parameter to `true` by using the following command:

```
1 xe host-param-set uuid=<host uuid> multipathing=true
```

4. If there are existing SRs on the server running in single path mode but that have multiple paths:
 - Migrate or suspend any running guests with virtual disks in affected the SRs
 - Replug the PBD of any affected SRs to reconnect them using multipathing:

```
1 xe pbd-plug uuid=<pbd_uuid>
```

5. Repeat these steps to enable multipathing on all hosts in the pool.

Ensure that you enable multipathing on all hosts in the pool. All cabling and, in the case of iSCSI, subnet configurations must match the corresponding NICs on each host.

Disable multipathing

You can disable multipathing in XenCenter or on the xe CLI.

To disable multipathing by using XenCenter

1. In the XenCenter **Resources** pane, right-click on the host and choose **Enter Maintenance Mode**.
2. Wait until the host reappears in the **Resources** pane with the maintenance mode icon (a blue square) before continuing.
3. On the **General** tab for the host, click **Properties** and then go to the **Multipathing** tab.
4. To disable multipathing, clear the **Enable multipathing on this server** check box.
5. Click **OK** to apply the new setting. There is a short delay while XenCenter saves the new storage configuration.

6. In the **Resources** pane, right-click on the host and choose **Exit Maintenance Mode**.
7. Repeat these steps to configure multipathing on all hosts in the pool.

To disable multipathing by using the xe CLI

1. Open a console on the XenServer host.
2. Unplug all PBDs on the host by using the following command:

```
1 xe pbd-unplug uuid=<pbd_uuid>
```

You can use the command `xe pbd-list` to find the UUID of the PBDs.

3. Set the value of the `multipathing` parameter to **false** by using the following command:

```
1 xe host-param-set uuid=<host uuid> multipathing=false
```

4. If there are existing SRs on the server running in single path mode but that have multiple paths:
 - Migrate or suspend any running guests with virtual disks in affected the SRs
 - Unplug and replug the PBD of any affected SRs to reconnect them using multipathing:

```
1 xe pbd-plug uuid=<pbd_uuid>
```

5. Repeat these steps to disable multipathing on all hosts in the pool.

Configure multipathing

To make additional temporary multipath configuration modifications, create files with the suffix `.conf` in the directory `/etc/multipath/conf.d`. Add the additional configuration in these files. Multipath searches the directory alphabetically for files ending in `.conf` and reads configuration information from them.

Do not edit the file `/etc/multipath.conf`. This file is overwritten by updates to XenServer. If you want to make permanent changes to your multipathing configuration, contact your storage vendor who can request our partner engineering team to have the changes permanently added.

iSCSI multipathing on a single subnet

You can configure XenServer to use iSCSI multipath with storage arrays that expose their targets and IQN(s) on a single subnet. For example, you can follow these steps to set up Dell EqualLogic PS and FS unified series storage arrays.

By default, XenServer establishes only one connection per iSCSI target. Hence, with the default configuration the recommendation is to use NIC bonding to achieve failover and load balancing. The configuration procedure outlined in this section describes an alternative configuration, where multiple iSCSI connections are established for a single subnet or target. NIC bonding is not required.

Note:

The following configuration is only supported for servers that are exclusively attached to storage arrays which expose their targets and IQN(s) through a single subnet. These storage arrays must be qualified for this procedure with XenServer.

To configure multipath:

1. Back up any data you want to protect.
2. In the XenCenter **Resources** pane, right-click on the host and choose **Enter Maintenance Mode**.
3. Wait until the host reappears in the **Resources** pane with the maintenance mode icon (a blue square) before continuing.
4. On the **General** tab for the host, click **Properties** and then go to the **Multipathing** tab.
5. To enable multipathing, select the **Enable multipathing on this server** check box.
6. Click **OK** to apply the new setting. There is a short delay while XenCenter saves the new storage configuration.
7. In the host console, configure two to four Open-iSCSI interfaces. Each iSCSI interface is used to establish a separate path. The following steps show the process for two interfaces:

- a) Configure two iSCSI interfaces, run the following commands:

```
1 iscsiadm -m iface --op new -I c_iface1
2 iscsiadm -m iface --op new -I c_iface2
```

Ensure that the interface names have the prefix `c_`. If the interfaces do not use this naming standard, they are ignored and instead the default interface is used.

Note:

This configuration leads to the default interface being used for all connections. This indicates that all connections are being established using a single interface.

- b) Bind the iSCSI interfaces to `xenbr1` and `xenbr2`, by using the following commands:

```
1 iscsiadm -m iface --op update -I c_iface1 -n iface.
  net_ifacename -v xenbr1
2 iscsiadm -m iface --op update -I c_iface2 -n iface.
  net_ifacename -v xenbr2
```

Note:

This configuration assumes that the network interfaces configured for the control domain (including xenbr1 and xenbr2) and xenbr0 are used for management. It also assumes that the NIC cards being used for the storage network are NIC1 and NIC2. If this is not the case, refer to your network topology to discover the network interfaces and NIC cards to use in these commands.

8. In the XenCenter **Resources** pane, right-click on the host and choose **Exit Maintenance Mode**. Do not resume your VMs yet.

9. In the host console, run the following commands to discover and log in to the sessions:

```
1 iscsiadm -m discovery -t st -p <IP of SAN>
2 iscsiadm -m node -L all
```

10. Delete the stale entries containing old session information by using the following commands:

```
1 cd /var/lib/iscsi/send_targets/<IP of SAN and port, use ls command
  to check that>
2 rm -rf <iqn of SAN target for that particular LUN>
3
4 cd /var/lib/iscsi/nodes/
5 rm -rf <entries for that particular SAN>
```

11. Detach the LUN and attach it again. You can do this in one of the following ways:

- After completing the preceding steps on all hosts in a pool, you can use XenCenter to detach and reattach the LUN for the entire pool.
- Alternatively, you can unplug and destroy the PBD for each host and then repair the SR.

- a) Run the following commands to unplug and destroy the PBD:

- i. Find the UUID of the SR:

```
1 xe sr-list
```

- ii. Get the list of PBDs associated with the SR:

```
1 xe pbd-list sr-uuid=<sr-uuid>
```

- iii. In the output of the previous command, look for the UUID of the PBD of the iSCSI Storage Repository with a mismatched SCSI ID.

- iv. Unplug and destroy the PBD you identified.

```
1 xe pbd-unplug uuid=<pbd-uuid>
2 xe pbd-destroy uuid=<pbd-uuid>
```

- b) Repair the storage in XenCenter.

12. You can now resume your VMs.

Storage read caching

December 17, 2025

Read caching improves a VM's disk performance as, after the initial read from external disk, data is cached within the host's free memory. It improves performance in situations where many VMs are cloned off a single base VM, as it drastically reduces the number of blocks read from disk. For example, in Citrix Virtual Desktops environment Machine Creation Services™ (MCS) environments.

The performance improvement can be seen whenever data is read from disk more than once, as it gets cached in memory. This change is most noticeable in the degradation of service that occurs during heavy I/O situations. For example, in the following situations:

- When a significant number of end users boot up within a very narrow time frame (boot storm)
- When a significant number of VMs are scheduled to run malware scans at the same time (antivirus storms).

Read caching is enabled by default when you have the appropriate license type.

Note:

Storage Read Caching is available for XenServer® Premium Edition customers.

Enable and disable read caching

For file-based SRs, such as NFS, EXT3/EXT4, SMB, and GFS2 SR types, read-caching is enabled by default. Read-caching is disabled for all other SRs.

You can choose to disable read caching on NFS, EXT3/EXT4, or SMB SRs. For GFS2 SRs, read caching cannot be disabled.

To disable read caching for a specific SR by using the xe CLI, run the following command:

```
1 xe sr-param-set uuid=sr-uuid other-config:o_direct=true
```

To disable read caching for a specific SR by using XenCenter®, go to the **Properties** dialog for the SR. In the **Read Caching** tab, you can select to enable or disable read caching.

For more information, see [Changing SR Properties](#).

Limitations

- Read caching is available only for NFS, EXT3/EXT4, SMB, and GFS2 SRs. It is not available for other SR types.
- Read caching only applies to read-only VDIs and VDI parents. These VDIs exist where VMs are created from 'Fast Clone' or disk snapshots. The greatest performance improvements can be seen when many VMs are cloned from a single 'golden' image.
- Performance improvements depend on the amount of free memory available in the host's Control Domain (dom0). Increasing the amount of dom0 memory allows more memory to be allocated to the read-cache. For information on how to configure dom0 memory, see [Change the amount of memory allocated to the control domain](#).
- When memory read caching is turned on, a cache miss causes I/O to become serialized. This can sometimes be more expensive than having read caching turned off, because with read caching turned off I/O can be parallelized. To reduce the impact of cache misses, increase the amount of available dom0 memory or disable read caching for the SR.

Comparison with IntelliCache™

IntelliCache and memory based read caching are to some regards complementary. IntelliCache not only caches on a different tier, but it also caches writes in addition to reads. IntelliCache caches reads from the network onto a local disk. In-memory read caching caches the reads from network or disk into host memory. The advantage of in-memory read caching, is that memory is still an order of magnitude faster than a solid-state disk (SSD). Performance in boot storms and other heavy I/O situations improves.

Both read-caching and IntelliCache can be enabled simultaneously. In this case, IntelliCache caches the reads from the network to a local disk. Reads from that local disk are cached in memory with read caching.

Set the read cache size

The read cache performance can be optimized, by giving more memory to XenServer's control domain (dom0).

Important:

Set the read cache size on ALL hosts in the pool individually for optimization. Any subsequent changes to the size of the read cache must also be set on all hosts in the pool.

On the XenServer host, open a local shell and log on as root.

To set the size of the read cache, run the following command:

```
1 /opt/xensource/libexec/xen-cmdline --set-xen dom0_mem=nnM,max:nnM
```

Set both the initial and maximum values to the same value. For example, to set dom0 memory to 20,480 MiB:

```
1 /opt/xensource/libexec/xen-cmdline --set-xen dom0_mem=20480M,max:20480M
```

Important:

Reboot all hosts after changing the read cache size.

How to view the current dom0 memory allocation?

To view the current dom0 memory settings, enter:

```
1 free -m
```

The output of `free -m` shows the current dom0 memory settings. The value may be less than expected due to various overheads. The example table below shows the output from a host with dom0 set to 2.6 GiB

1				Total		Used		Free		Shared		Buffer/cache	
2		Available											
3		Mem:		2450		339		1556		9		554	
4		Swap:		1023		0		1023					

What Range of Values Can be Used? As the XenServer Control Domain (dom0) is 64-bit, large values can be used, for example 32,768 MiB. However, we recommend that you **do not reduce the dom0 memory below 1 GiB**.

XenCenter display notes

The entire host’s memory can be considered to comprise the Xen® hypervisor, dom0, VMs, and free memory. Even though dom0 and VM memory is usually of a fixed size, the Xen hypervisor uses a variable amount of memory. The amount of memory used depends on various factors. These factors include the number of VMs running on the host at any time and how those VMs are configured. It is not possible to limit the amount of memory that Xen uses. Limiting the amount of memory can cause Xen to run out of memory and prevent new VMs from starting, even when the host had free memory.

To view the memory allocated to a host, in XenCenter select the host, and then click the **Memory** tab.

The XenServer field displays the *sum* of the memory allocated to dom0 *and* Xen memory. Therefore, the amount of memory displayed might be higher than specified by the administrator. The memory size can vary when starting and stopping VMs, even when the administrator has set a fixed size for dom0.

Graphics overview

June 24, 2026

This section provides an overview of the virtual delivery of 3D professional graphics applications and workstations in XenServer. The offerings include GPU Pass-through (for NVIDIA, AMD, and Intel GPUs) and hardware-based GPU sharing with NVIDIA vGPU.

Graphics Virtualization is available for XenServer Premium Edition customers. To learn more about XenServer editions, and to find out how to upgrade, visit the [XenServer website](#). For more information, see [Licensing](#).

GPU pass-through

In a virtualized system, most of the physical system components are shared. These components are represented as multiple virtual instances to multiple clients by the hypervisor. A pass-through GPU is not abstracted at all, but remains one physical device. Each hosted virtual machine (VM) gets its own dedicated GPU, eliminating the software abstraction and the performance penalty that goes with it.

XenServer allows you to assign a physical GPU (in the XenServer host) to a Windows or Linux VM running on the same host. This GPU pass-through feature is intended for graphics power users, such as CAD designers.

Shared GPU (vGPU)

Shared GPU (vGPU) allows one physical GPU to be used by multiple VMs concurrently. Because a portion of a physical GPU is used, performance is greater than emulated graphics, and there is no need for one card per VM. This feature enables resource optimization, boosting the performance of the VM. The graphics commands of each virtual machine are passed directly to the GPU, without translation by the hypervisor.

Multiple shared GPU (vGPU)

Multiple vGPU enables multiple virtual GPUs to be used concurrently by a single VM. Only certain vGPU profiles can be used and all vGPUs attached to a single VM must be of the same type. These additional vGPUs can be used to perform computational processing. For more information about the number of vGPUs supported for a single VM, see [Configuration Limits](#).

This feature is only available for NVIDIA GPUs. For more information about the physical GPUs that support the multiple vGPU feature, see the NVIDIA documentation.

vGPU vendor support

The following table lists guest support for the GPU pass-through, shared GPU (vGPU), and multiple shared GPU (vGPU) features:

	GPU pass-through for Windows VMs	GPU pass-through for Linux VMs	Shared GPU (vGPU) for Windows VMs	Shared GPU (vGPU) for Linux VMs	Multiple shared GPU (vGPU) for Windows VMs	Multiple shared GPU (vGPU) for Linux VMs
AMD	YES					
Intel	YES					
NVIDIA	YES	YES	YES (preview)	YES (preview)	YES (preview, see note)	YES (preview, see note)

Note:

- NVIDIA vGPU support for XenServer 9 is currently in preview, and not fully supported.
- Only some of the guest operating systems support multiple vGPU. For more information, see [Guest support and constraints](#).
- Only some of the guest operating systems support vGPU live migration. For more information, see vGPU live migration vendor support.

You might need a vendor subscription or a license depending on the graphics card used.

vGPU live migration

vGPU live migration enables a VM that uses a virtual GPU to perform live migration, storage live migration, or VM suspend. VMs with vGPU live migration capabilities can be migrated to avoid downtime.

vGPU live migration also enables you to perform rolling pool upgrades on pools that host vGPU-enabled VMs. For more information, see [Rolling pool upgrades](#).

To use vGPU live migration or VM suspend, your VM must run on a graphics card that supports this feature. Your VM must also have the supported drivers from the GPU vendor installed.

Warning:

The size of the GPU state in the NVIDIA driver can cause a downtime of 5 seconds or more during vGPU live migration.

The following restrictions apply when using vGPU live migration:

- Live migration is not compatible with GPU Pass-through.
- VMs must have the appropriate vGPU drivers installed to be supported with any vGPU live migration features. The in-guest drivers must be installed for all guests using the vGPU feature.
- Reboot and shutdown operations on a VM are not supported while a migration is in progress. These operations can cause the migration to fail.
- Linux VMs are not supported with any vGPU live migration features.
- Live migration by the Workload Balancing appliance is not supported for vGPU-enabled VMs. The Workload Balancing appliance cannot do capacity planning for VMs that have a vGPU attached.
- After migrating a VM using vGPU live migration, the guest VNC console might become corrupted. Use ICA®, RDP, or another network-based method for accessing VMs after a vGPU live migration has been performed.
- VDI migration uses live migration, therefore requires enough vGPU space on the host to make a copy of the vGPU instance on the host. If the physical GPUs are fully used, VDI migration might not be possible.

vGPU live migration vendor support

The following table lists support for vGPU live migration:

GPU pass-through for Windows VMs		GPU pass-through for Linux VMs		Shared GPU (vGPU) for Windows VMs		Shared GPU (vGPU) for Linux VMs		Multiple shared GPU (vGPU) for Windows VMs	Multiple shared GPU (vGPU) for Linux VMs
NVIDIA				YES				YES	

For more information about the graphics cards that support this feature, see the vendor-specific sections of this guide. Customers might need a vendor subscription or a license depending on the graphics card used.

Guest support and constraints

XenServer supports the following guest operating systems for virtual GPUs.

NVIDIA vGPU

Operating systems marked with an asterisk (*) also support multiple vGPU.

Windows guests:

- Windows 10 (64-bit) *
- Windows 11 *
- Windows Server 2016 *
- Windows Server 2019 *
- Windows Server 2022 *
- Windows Server 2025 *

Linux guests:

- RHEL 8 *
- RHEL 9 *
- RHEL 10 (preview) *
- CentOS Stream 9
- CentOS Stream 10
- Ubuntu 22.04 *
- Ubuntu 24.04 *
- Rocky Linux 8 *
- Rocky Linux 9 *

Constraints

- VMs with a virtual GPU are not supported with Dynamic Memory Control.
- XenServer automatically detects and groups identical physical GPUs across hosts in the same pool. If assigned to a group of GPUs, a VM can be started on any host in the pool that has an available GPU in the group.

- All graphics solutions (NVIDIA vGPU and vGPU pass-through) can be used in an environment that uses high availability. However, VMs that use these graphics solutions cannot be protected with high availability. These VMs can be restarted on a best-effort basis while there are hosts with the appropriate free resources.

Prepare host for graphics

June 22, 2026

This section provides step-by-step instructions on how to prepare XenServer® for supported graphical virtualization technologies. The offerings include NVIDIA vGPU.

NVIDIA vGPU

NVIDIA vGPU enables multiple Virtual Machines (VM) to have simultaneous, direct access to a single physical GPU. It uses NVIDIA graphics drivers deployed on non-virtualized Operating Systems. NVIDIA physical GPUs can support multiple virtual GPU devices (vGPUs). To provide this support, the physical GPU must be under the control of NVIDIA Virtual GPU Manager running in XenServer Control Domain (dom0). The vGPUs can be assigned directly to VMs.

VMs use virtual GPUs like a physical GPU that the hypervisor has passed through. An NVIDIA driver loaded in the VM provides direct access to the GPU for performance critical fast paths. It also provides a paravirtualized interface to the NVIDIA Virtual GPU Manager.

Important:

To ensure that you always have the latest security and functional fixes, ensure that you install the latest NVIDIA vGPU software package for XenServer (consisting of the NVIDIA Virtual GPU Manager for XenServer and NVIDIA drivers) and keep it updated to the latest version provided by NVIDIA. For more information, see [the NVIDIA documentation](#).

The latest NVIDIA drivers are available from the [NVIDIA Application Hub](#).

NVIDIA vGPU is compatible with the HDX 3D Pro feature of Citrix Virtual Apps and Desktops or Citrix DaaS. For more information, see [HDX 3D Pro](#).

For more information about NVIDIA vGPUs, see the [NVIDIA Documentation Hub](#).

Licensing note

NVIDIA vGPU is available for XenServer Premium Edition customers. To learn more about XenServer editions, and to find out how to upgrade, visit the [XenServer website](#). For more information, see [Li-](#)

censing.

Depending on the NVIDIA graphics card used, you might need NVIDIA subscription or a license.

For information on licensing NVIDIA cards, see the [NVIDIA website](#).

Available NVIDIA vGPU types

NVIDIA GRID cards contain multiple Graphics Processing Units (GPU). For example, TESLA M10 cards contain four GM107GL GPUs, and TESLA M60 cards contain two GM204GL GPUs. Each physical GPU can host several different types of virtual GPU (vGPU). vGPU types have a fixed amount of frame buffer, number of supported display heads and maximum resolutions, and are targeted at different classes of workload.

For a list of the most recently supported NVIDIA cards, see the [Hardware Compatibility List](#) and the [NVIDIA product information](#).

Note:

The vGPUs hosted on a physical GPU at the same time **must all be of the same type**. However, there is no corresponding restriction for physical GPUs on the same card. This restriction is automatic and can cause unexpected capacity planning issues.

NVIDIA vGPU system requirements

- NVIDIA GRID card:
 - For a list of the most recently supported NVIDIA cards, see the [Hardware Compatibility List](#) and the [NVIDIA product information](#).
- Depending on the NVIDIA graphics card used, you might need an NVIDIA subscription or a license. For more information, see the [NVIDIA product information](#).
- Depending on the NVIDIA graphics card, you might need to ensure that the card is set to the correct mode. For more information, see the [NVIDIA documentation](#).
- XenServer Premium Edition.
- A host capable of hosting XenServer and the supported NVIDIA cards.
- NVIDIA vGPU software package for XenServer, consisting of the NVIDIA Virtual GPU Manager for XenServer, and NVIDIA drivers.

Note:

Review the NVIDIA Virtual GPU software documentation available from the [NVIDIA website](#). Register with NVIDIA to access these components.

- To run Citrix Virtual Desktops with VMs running NVIDIA vGPU, you also need: Citrix Virtual Desktops 7.6 or later, full installation.
- For NVIDIA Ampere vGPUs and all future generations, you must enable SR-IOV in your system firmware.

vGPU live migration

XenServer enables the use of live migration, storage live migration, and the ability to suspend and resume for NVIDIA vGPU-enabled VMs.

To use the vGPU live migration, storage live migration, or Suspend features, satisfy the following requirements:

- An NVIDIA GRID card, Maxwell family or later.
- An NVIDIA Virtual GPU Manager for XenServer with live migration enabled. For more information, see the NVIDIA Documentation.
- A Windows VM that has NVIDIA live migration-enabled vGPU drivers installed.

vGPU live migration enables the use of live migration within a pool, live migration between pools, storage live migration, and Suspend/Resume of vGPU-enabled VMs.

Preparation overview

1. Install XenServer
2. Install the NVIDIA Virtual GPU Manager for XenServer
3. Restart the XenServer host

Installation on XenServer

XenServer is available for download from the [XenServer Downloads](#) page.

Install the following:

- **XenServer Base Installation ISO**
- **XenCenter® Windows Management Console**

For more information, see [Install](#).

Licensing note

vGPU is available for XenServer Premium Edition customers. To learn more about XenServer editions, and to find out how to upgrade, visit the [XenServer website](#). For more information, see [Licensing](#).

Depending on the NVIDIA graphics card used, you might need NVIDIA subscription or a license. For more information, see [NVIDIA product information](#).

For information about licensing NVIDIA cards, see the [NVIDIA website](#).

Install the NVIDIA vGPU Manager for XenServer

Install the NVIDIA Virtual GPU software that is available from [NVIDIA](#). The NVIDIA Virtual GPU software consists of:

- NVIDIA Virtual GPU Manager
- Windows Display Driver (The Windows display driver depends on the Windows version)

The **NVIDIA Virtual GPU Manager** runs in the XenServer Control Domain (dom0). It is provided as a supplemental pack. For more information about installation, see the [NVIDIA virtual GPU Software documentation](#).

Important:

Do not install the NVIDIA Virtual GPU Manager supplemental pack as part of a fresh installation of a XenServer host. After you have installed your hosts, you can install NVIDIA Virtual GPU Manager.

The update can be installed using the following method:

1. Use XenCenter (**Tools > Install Update > Select update or supplemental pack from disk**), or the xe CLI command `xe-install-supplemental-pack`.
2. Restart the XenServer host
3. After you restart the XenServer host, you can verify that the software has been installed and loaded correctly by checking the NVIDIA kernel driver in dom0:

```
1 [root@xenserver ~]#lsmod |grep nvidia
2      nvidia                8152994  0
```

4. Verify that the NVIDIA kernel driver can successfully communicate with the NVIDIA physical GPUs in your host. Run the `nvidia-smi` command to produce a listing of the GPUs in your platform similar to:

```
1 [root@xenserver ~]# nvidia-smi
2
3      Thu Jan 26 13:48:50 2017
```

4	+-----+									
5	NVIDIA-SMI 367.64 Driver Version: 367.64									
6	-----+									
7	GPU Name Persistence-M Bus-Id Disp.A Volatile Uncorr.									
	ECC									
8	Fan Temp Perf Pwr:Usage/Cap Memory-Usage GPU-Util									
	Compute M.									
9	=====+									
10	0 Tesla M60 On 0000:05:00.0 Off Off									
11	N/A 33C P8 24W / 150W 7249MiB / 8191MiB 0%									
	Default									
12	+-----+									
13	1 Tesla M60 On 0000:09:00.0 Off Off									
14	N/A 36C P8 24W / 150W 7249MiB / 8191MiB 0%									
	Default									
15	+-----+									
16	2 Tesla M60 On 0000:85:00.0 Off Off									
17	N/A 36C P8 23W / 150W 19MiB / 8191MiB 0%									
	Default									
18	+-----+									
19	3 Tesla M60 On 0000:89:00.0 Off Off									
20	N/A 37C P8 23W / 150W 14MiB / 8191MiB 0%									
	Default									
21	+-----+									
22	+-----+									
23	Processes: GPU Memory									
24	GPU PID Type Process name Usage									
25	=====									
26	No running compute processes found									
27	+-----+									

Note:

When using NVIDIA vGPU with XenServer servers that have more than 768 GB of RAM, add the parameter `dom0-iommu=passthrough` to the Xen command line:

- a) Run the following command in the control domain (Dom0):
`/opt/xensource/libexec/xen-cmdline --set-xen dom0-iommu=passthrough`
- b) Restart the host.

Intel GPU pass-through


XenServer supports the GPU pass-through feature for Windows VMs using an Intel integrated GPU device.

- For more information on Windows versions supported with Intel GPU pass-through, see [Graphics](#).
- For more information on supported hardware, see the [Hardware Compatibility List](#).

When using Intel GPU on Intel servers, the XenServer server's Control Domain (dom0) has access to the integrated GPU device. In such cases, the GPU is available for pass-through. To use the Intel GPU Pass-through feature on Intel servers, disable the connection between dom0 and the GPU before passing through the GPU to the VM.

To disable this connection, complete the following steps:

1. On the **Resources** pane, choose the XenServer host.
2. On the **General** tab, click **Properties**, and in the left pane, click **GPU**.
3. In the **Integrated GPU passthrough** section, select **This server will not use the integrated GPU**.



This step disables the connection between dom0 and the Intel integrated GPU device.

4. Click **OK**.
5. Restart the XenServer host for the changes to take effect.

The Intel GPU is now visible on the GPU type list during new VM creation, and on the VM's **Properties** tab.

Note:

The XenServer host's external console output (for example, VGA, HDMI, DP) will not be available after disabling the connection between dom0 and the GPU.

Create vGPU enabled VMs


May 13, 2026

This section provides step-by-step instructions on how to create a virtual GPU or GPU pass-through enabled VM.

Note:

If you are using the Intel GPU Pass-through feature, first see the section *Enabling Intel GPU Pass-through* for more configuration, and then complete the following steps.

1. Create a VM using XenCenter. Select the host on the Resources pane and then select **New VM** on the VM menu.
2. Follow the instructions on the **New VM** configuration and select the **Installation Media**, **Home Server**, and **CPU & Memory**.
3. GPU-enabled hosts display a **GPU** configuration page:



4. Click **Add**. From the **GPU Type** list, select either **Pass-through whole GPU**, or a virtual GPU type.

Unavailable virtual GPU types are grayed-out.

If you want to assign multiple vGPUs to your VM, ensure that you select a vGPU type that supports multiple vGPU. Repeat this step to add more vGPUs of the same type.

5. Click **Next** to configure **Storage** and then **Networking**.
6. After you complete your configuration, click **Create Now**.

Install the XenServer® VM Tools

Without the optimized networking and storage drivers provided by the XenServer VM Tools, remote graphics applications running on NVIDIA vGPU do **not** deliver maximum performance.

- If your VM is a Windows VM, you must install the XenServer VM Tools for Windows on your VM. For more information, see [Install XenServer VM Tools for Windows](#).
- If your VM is a Linux VM, you can install the XenServer VM Tools for Linux on your VM. For more information, see [Install XenServer VM Tools for Linux](#).

Install the in-guest drivers

When viewing the VM console in XenCenter, the VM typically boots to the desktop in VGA mode with 800 x 600 resolution. The standard Windows screen resolution controls can be used to increase the resolution to other standard resolutions. (**Control Panel > Display > Screen Resolution**)

Note:

When using GPU pass-through, we recommend that you install the in-guest drivers through RDP or VNC over the network. That is, not through XenCenter.


To ensure that you always have the latest security and functional fixes, ensure that you always take the latest updates to your in-guest drivers.

Install the NVIDIA drivers

To enable vGPU operation (as for a physical NVIDIA GPU), install NVIDIA drivers into the VM.

The following section provides an overview of the procedure. For detailed instructions, see the [NVIDIA virtual GPU Software documentation](#).

1. Start the VM. In the **Resources** pane, right-click on the VM, and click **Start**.
During this start process, XenServer dynamically allocates a vGPU to the VM.
2. Follow the Windows operating system installation screens.
3. After the operating system installation completes, restart the VM.
4. Install the appropriate driver for the GPU inside the guest. The following example shows the specific case for in guest installation of the NVIDIA GRID drivers.
5. Copy the 64-bit NVIDIA Windows driver package to the VM, open the zip file, and run setup.exe.
6. Follow the installer steps to install the driver.
7. After the driver installation has completed, you might be prompted to reboot the VM. Select **Restart Now** to restart the VM immediately, alternatively, exit the installer package, and restart the VM when ready. When the VM starts, it boots to a Windows desktop.
8. To verify that the NVIDIA driver is running, right-click on the desktop and select **NVIDIA Control Panel**.
9. In the NVIDIA Control Panel, select **System Information**. This interface shows the GPU Type in use by the VM, its features, and the NVIDIA driver version in use:

**Note:**

Depending on the NVIDIA graphics card used, you might need an NVIDIA subscription or a

license. For more information, see the [NVIDIA product information](#).

The VM is now ready to run the full range of DirectX and OpenGL graphics applications supported by the GPU.

For Linux VMs

In your Linux VM, install the driver as instructed by the NVIDIA User Guides.

If your Linux VM boots in UEFI Secure Boot mode, you might be required to take additional steps to sign the driver. For more information, see [Install third-party drivers on your Secure Boot Linux VM](#).

Memory usage

June 30, 2026

Two components contribute to the memory footprint of the XenServer® host. First, the memory consumed by the Xen hypervisor itself. Second, there is the memory consumed by the *Control Domain* of the host. Also known as ‘Domain0’, or ‘dom0’, the control domain is a secure, privileged Linux VM that runs the XenServer management toolstack (XAPI). Besides providing XenServer management functions, the control domain also runs the driver stack that provides user created VM access to physical devices.

Control domain memory

The amount of memory allocated to the control domain is adjusted automatically and is based on the amount of physical memory on the physical host. By default, XenServer allocates **1 GiB plus 5% of the total physical memory** to the control domain, up to an initial maximum of 8 GiB.

Note:

The amount reported in the XenServer section in XenCenter® includes the memory used by the control domain (dom0), the Xen hypervisor itself, and the crash kernel. Therefore, the amount of memory reported in XenCenter can exceed these values. The amount of memory used by the hypervisor is larger for hosts using more memory.

Change the amount of memory allocated to the control domain

You can change the amount of memory allocated to dom0 by using XenCenter or by using the command line. If you increase the amount of memory allocated to the control domain beyond the amount

allocated by default, this action results in less memory being available to VMs.

You might need to increase the amount of memory assigned to the control domain of a XenServer host in the following cases:

- You are running many VMs on the host
- You are using PVS-Accelerator
- You are using read caching

Important:

If you are using a GFS2 SR and any of these cases also applies to your environment, you must increase the amount of control domain memory. Insufficient control domain memory can cause network instability, which can cause problems for clustered pools with GFS2 SRs.

The amount of memory to allocate to the control domain depends on your environment and the requirements of your VMs.

You can monitor the following metrics to judge whether the amount of control domain memory is appropriate for your environment and what effects any changes you make have:

- **Swap activity:** If the control domain is swapping, increase the control domain memory.
- **Tapdisk mode:** You can monitor whether your tapdisks are in low-memory mode from within the XenCenter **Performance** tab for the host. Select **Actions > New Graph** and choose the **Tapdisks in low memory mode** graph. If a tapdisk is in low-memory mode, increase the control domain memory.
- **Pagecache pressure:** Use the `top` command to monitor the `buff/cache` metric. If this number becomes too low, you might want to increase the control domain memory.

Changing the dom0 memory by using XenCenter

For information about changing the dom0 memory by using XenCenter, see [Changing the Control Domain Memory](#) in the XenCenter documentation.

Note:

You cannot use XenCenter to reduce dom0 memory below the value that was initially set during XenServer installation. To make this change you must use the command line.

Changing the dom0 memory by using the command line**Note:**

On hosts with smaller memory (less than 16 GiB), you might want to reduce the memory allocated to the Control Domain to lower than the installation default value. You can use the com-

mand line to make this change. However, we recommend that you **do not reduce the dom0 memory below 1 GiB** and that you do this operation under the guidance of the Support Team.

1. On the XenServer host, open a local shell and log on as root.
2. Type the following:

```
1 xe host-list uuid=host_uuid params=control-domain-uuid
2 xe vm-memory-set uuid=<control domain uuid> memory=<nn>M
```

Where **<nn>** represents the amount of memory, in MiB, to be allocated to dom0.

3. Restart the XenServer host using XenCenter or the **reboot** command on the XenServer console.

When the host restarts, on the XenServer console, run the **free** command to verify the new memory settings.

How much memory is available to VMs?

To find out how much host memory is available to be assigned to VMs, find the value of the free memory of the host by running **memory-free**. Then type the command **vm-compute-maximum-memory** to get the actual amount of free memory that can be allocated to the VM. For example:

```
1 xe host-list uuid=host_uuid params=memory-free
2 xe vm-compute-maximum-memory vm=vm_name total=host_memory_free_value
```

Monitor and manage your deployment

December 17, 2025

XenServer provides detailed monitoring of performance metrics. These metrics include CPU, memory, disk, network, C-state/P-state information, and storage. Where appropriate, these metrics are available on a per host and a per VM basis. These metrics are available directly, or can be accessed and viewed graphically in XenCenter or other third-party applications.

XenServer also provides system and performance alerts. Alerts are notifications that occur in response to selected system events. These notifications also occur when one of the following values goes over a specified threshold on a managed host, VM, or storage repository: CPU usage, network usage, memory usage, control domain memory usage, storage throughput, or VM disk usage. You can configure the alerts by using the xe CLI or by using XenCenter. To create notifications based on any of the available Host or VM performance metrics see [Performance alerts](#).

Monitor XenServer® performance

Customers can monitor the performance of their XenServer hosts and Virtual Machines (VMs) using the metrics exposed through Round Robin Databases (RRDs). These metrics can be queried over HTTP or through the RRD2CSV tool. In addition, XenCenter uses this data to produce system performance graphs. For more information, see [Analyze and visualize metrics](#).

Analyze and visualize metrics

The Performance tab in XenCenter provides real time monitoring of performance statistics across resource pools in addition to graphical trending of virtual and physical machine performance. Graphs showing CPU, memory, network, and disk I/O are included on the Performance tab by default. You can add more metrics, change the appearance of the existing graphs or create extra ones. For more information, see *Configuring metrics* in the following section.

- You can view up to 12 months of performance data and zoom in to take a closer look at activity spikes.
- XenCenter can generate performance alerts when CPU, memory, network I/O, storage I/O, or disk I/O usage exceed a specified threshold on a host, VM, or SR. For more information, see *Alerts* in the following section.

Note:

Install the XenServer VM Tools to see full VM performance data.

Configure performance graphs To add a graph:

1. On the **Performance** tab, click **Actions** and then **New Graph**. The New Graph dialog box is displayed.
2. In the **Name** field, enter a name for the graph.
3. From the list of **Datasources**, select the check boxes for the datasources you want to include in the graph.
4. Click **Save**.

To edit an existing graph:

1. Navigate to the **Performance** tab, and select the graph that you would like to modify.
2. Right-click on the graph and select **Actions**, or click the **Actions** button. Then select **Edit Graph**.
3. On the graph details window, make the necessary changes, and click **OK**.

Configure the graph type Data on the performance graphs can be displayed as lines or as areas. To change the graph type:

1. On the **Tools** menu, click **Options** and select **Graphs**.
2. To view performance data as a line graph, click the **Line graph** option.
3. To view performance data as an area graph, click the **Area graph** option.
4. Click **OK** to save your changes.

Comprehensive details for configuring and viewing XenCenter performance graphs can be found in the XenCenter documentation in the section [Monitoring System Performance](#).

Configure metrics

Note:

C-states and P-states are power management features of some processors. The range of states available depends on the physical capabilities of the host, as well power management configuration.

Both host and VM commands return the following:

- A full description of the data source
- The units applied to the metric
- The range of possible values that may be used

For example:

```
1  name_label: cpu0-C1
2  name_description: Proportion of time CPU 0 spent in C-state 1
3  enabled: true
4  standard: true
5  min: 0.000
6  max: 1.000
7  units: Percent
```

Enable a specific metric Most metrics are enabled and collected by default, to enable those metrics that are not, enter the following:

```
1  xe host-data-source-record data-source=metric name host=hostname
```

Disable a specific metric You might not want to collect certain metrics regularly. To disable a previously enabled metric, enter the following:

```
1 xe host-data-source-forget data-source=metric name host=hostname
```

Display a list of currently enabled host metrics To list the host metrics currently being collected, enter the following:

```
1 xe host-data-source-list host=hostname
```

Display a list of currently enabled VM metrics To host the VM metrics currently being collected, enter the following:

```
1 xe vm-data-source-list vm=vm_name
```

Use RRDs

XenServer uses RRDs to store performance metrics. These RRDs consist of multiple Round Robin Archives (RRAs) in a fixed size database.

Each archive in the database samples its particular metric on a specified granularity:

- Every 5 seconds for 10 minutes
- Every minute for the past two hours
- Every hour for the past week
- Every day for the past year

The sampling that takes place every five seconds records actual data points, however the following RRAs use Consolidation Functions instead. The consolidation functions supported by XenServer are:

- AVERAGE
- MIN
- MAX

RRDs exist for individual VMs (including dom0) and the XenServer host. VM RRDs are stored on the host on which they run, or the pool coordinator when not running. Therefore the location of a VM must be known to retrieve the associated performance data.

For detailed information on how to use XenServer RRDs, see the [XenServer Software Development Kit Guide](#).

Analyze RRDs using HTTP

You can download RRDs over HTTP from the XenServer host specified using the HTTP handler registered at `/host_rrd` or `/vm_rrd`. Both addresses require authentication either by HTTP authentication, or by providing a valid management API session references as a query argument. For example:

Download a Host RRD.

```
1 wget http://server/host_rrd?session_id=OpaqueRef:SESSION_HANDLE>
```

Download a VM RRD.

```
1 wget http://server/vm_rrd?session_id=OpaqueRef:SESSION_HANDLE>&uuid=VM
  UUID>
```

Both of these calls download XML in a format that can be parsed directly.

Analyze RRDs using rrd2csv

In addition to viewing performance metrics in XenCenter®, the `rrd2csv` tool logs RRDs to Comma Separated Value (CSV) format. Man and help pages are provided. To display the `rrd2csv` tool man or help pages, run the following command:

```
1 man rrd2csv
```

Or

```
1 rrd2csv --help
```

Note:

Where multiple options are used, supply them individually. For example: to return both the UUID and the name-label associated with a VM or a host, call `rrd2csv` as shown below:

```
rrd2csv -u -n
```

The UUID returned is unique and suitable as a primary key, however the name-label of an entity might not necessarily be unique.

The man page (`rrd2csv --help`) is the definitive help text of the tool.

Available metrics

The following tables list all of the available host and VM metrics.

Notes:

- Latency over a period is defined as the average latency of operations during that period.
- The availability and utility of certain metrics are SR and CPU dependent.

Available host metrics

Metric Name	Description	Condition	XenCenter Name
<code>avgqu_sz_<sr-uuid-short></code>	Average I/O queue size (requests).	At least one plugged VBD in SR <code><sr-uuid-short></code> on the host	<code>sr-uuid-short</code> Queue Size
<code>cpu<cpu>-C<cstate></code>	Time CPU <code>cpu</code> spent in C-state <code>cstate</code> in milliseconds.	C-state exists on CPU	CPU <code>cpu</code> C-state <code>cstate</code>
<code>cpu<cpu>-P<pstate></code>	Time CPU <code>cpu</code> spent in P-state <code>pstate</code> in milliseconds.	P-state exists on CPU	CPU <code>cpu</code> P-state <code>pstate</code>
<code>cpu<cpu></code>	Utilization of physical CPU <code>cpu</code> (fraction). Enabled by default.	CPU <code>cpu</code> exists	CPU <code>cpu</code>
<code>cpu_avg</code>	Mean utilization of physical CPUs (fraction). Enabled by default.	None	Average CPU
<code>DCMI-power-reading</code>	Host power usage reported by IPMI DCMI.	Host hardware supports IPMI DCMI power reading	DCMI-power-reading
<code>hostload</code>	Host load per physical CPU, where load refers to the number of vCPU(s) in a running or runnable state.	None	Host CPU Load
<code>inflight_<sr-uuid-short></code>	Number of I/O requests currently in flight. Enabled by default.	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Inflight Requests

Metric Name	Description	Condition	XenCenter Name
<code>io_throughput_read_<sr-uuidshort></code>	Data read from SR (MiB/s).	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Read Throughput
<code>io_throughput_write_<sr-uuidshort></code>	Data written to the SR (MiB/s).	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Write Throughput
<code>io_throughput_total_<sr-uuidshort></code>	All SR I/O (MiB/s).	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Total Throughput
<code>iops_read_<sr-uuid-short></code>	Read requests per second.	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Read IOPS
<code>iops_write_<sr-uuid-short></code>	Write requests per second.	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Write IOPS
<code>iops_total_<sr-uuid-short></code>	I/O requests per second.	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Total IOPS
<code>iowait_<sr-uuid-short></code>	Percentage of the time waiting for I/O.	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> IO Wait
<code>latency_<sr-uuid-short></code>	Average I/O latency (milliseconds).	At least one plugged VBD in SR <code>sr</code> on the host	<code>sr</code> Latency
<code>loadavg</code>	Domain0 load average. Enabled by default	None	Control Domain Load
<code>memory_free_kib</code>	Total amount of free memory (KiB). Enabled by default.	None	<i>Not present in XenCenter. Replaced by Used Memory.</i>
<i>Not reported by the toolstack. Calculated by XenCenter.</i>	Total amount of used memory (KiB). Enabled by default.	None	Used Memory
<code>memory_reclaimed</code>	Host memory reclaimed by squeeze (B).	None	Reclaimed Memory

Metric Name	Description	Condition	XenCenter Name
<code>memory_reclaimed_mb</code>	Host memory available to reclaim with squeeze (B).	None	Potential Reclaimed Memory
<code>memory_total_kib</code>	Total amount of memory (KiB) in the host. Enabled by default.	None	Total Memory
<code>network/latency</code>	Interval in seconds between the last two heartbeats transmitted from the local host to all online hosts. Disabled by default.	HA Enabled	Network Latency
<code>statefile/<vdi_uuid>/latency</code>	Turn-around time in seconds of the latest State-File access from the local host. Disabled by default.	HA Enabled	HA State File Latency
<code>pif_<pif>_rx</code>	Bytes per second received on physical interface <code>pif</code> . Enabled by default.	PIF exists	XenCenter- <code>pifname</code> Receive (see note)
<code>pif_<pif>_tx</code>	Bytes per second sent on physical interface <code>pif</code> . Enabled by default.	PIF exists	XenCenter- <code>pifname</code> Send (see note)
<code>pif_<pif>_rx_errors</code>	Receive errors per second on physical interface <code>pif</code> . Disabled by default.	PIF exists	XenCenter- <code>pifname</code> Receive Errors (see note)
<code>pif_<pif>_tx_errors</code>	Transmit errors per second on physical interface <code>pif</code> . Disabled by default.	PIF exists	XenCenter- <code>pifname</code> Send Errors (see note)

Metric Name	Description	Condition	XenCenter Name
<code>pif_aggr_rx</code>	Bytes per second received on all physical interfaces. Enabled by default.	None	Total NIC Receive
<code>pif_aggr_tx</code>	Bytes per second sent on all physical interfaces. Enabled by default.	None	Total NIC Send
<code>pvsaccelerator_eviction_rate</code>	Bytes per second evicted from the cache	PVSAccelerator Enabled	PVS-Accelerator eviction rate
<code>pvsaccelerator_read_rate</code>	Reads per second served from the cache	PVSAccelerator Enabled	PVS-Accelerator hit rate
<code>pvsaccelerator_read_miss_rate</code>	Reads per second that cannot be served from the cache	PVSAccelerator Enabled	PVS-Accelerator miss rate
<code>pvsaccelerator_traffic_clients</code>	Bytes per second sent by cached PVS clients	PVSAccelerator Enabled	PVS-Accelerator observed network traffic from clients
<code>pvsaccelerator_traffic_servers</code>	Bytes per second sent by cached PVS servers	PVSAccelerator Enabled	PVS-Accelerator observed network traffic from servers
<code>pvsaccelerator_read_rate</code>	Reads per second observed by the cache	PVSAccelerator Enabled	PVS-Accelerator observed read rate
<code>pvsaccelerator_traffic_saved</code>	Bytes per second sent by PVSAccelerator instead of the PVS server	PVSAccelerator Enabled	PVS-Accelerator saved network traffic
<code>pvsaccelerator_space_utilization</code>	Percentage of space used by PVSAccelerator on this host, compared to the total size of the cache storage	PVSAccelerator Enabled	PVS-Accelerator space utilization
<code>running_vcpus</code>	The total number of running vCPUs	None	Number of running vCPUs

Metric Name	Description	Condition	XenCenter Name
running_domains	The total number of running domains including dom0 (the control domain of the host)	None	Number of running domains
sr_<sr>_cache_size	Size in bytes of the IntelliCache SR. Enabled by default.	IntelliCache Enabled	IntelliCache Cache Size
sr_<sr>_cache_hits	Cache hits per second. Enabled by default.	IntelliCache Enabled	IntelliCache Cache Hits
sr_<sr>_cache_misses	Cache misses per second. Enabled by default.	IntelliCache Enabled	IntelliCache Cache Misses
xapi_allocation_kib	Memory (KiB) allocation done by the XAPI daemon. Enabled by default.	None	Agent Memory Allocation
xapi_free_memory_kib	Free memory (KiB) available to the XAPI daemon. Enabled by default.	None	Agent Memory Free
xapi_healthcheck/latency	Turn-around time in seconds of the latest XAPI status monitoring call on the local host. Disabled by default.	High availability Enabled	XenServer High Availability Latency
xapi_live_memory_kib	Live memory (KiB) used by XAPI daemon. Enabled by default.	None	Agent Memory Live
xapi_memory_usage_total	Total memory (KiB) allocated used by XAPI daemon. Enabled by default.	None	Agent Memory Usage

Available VM metrics

Metric Name	Description	Condition	XenCenter Name
<code>cpu<cpu></code>	Utilization of vCPU <code>cpu</code> (fraction). Enabled by default	vCPU <code>cpu</code> exists	CPU
<code>cpu_usage</code>	Domain CPU usage	None	<code>cpu_usage</code>
<code>memory</code>	Memory currently allocated to VM (Bytes). Enabled by default	None	Total Memory
<code>memory_target</code>	Target of VM balloon driver (Bytes). Enabled by default	None	Memory target
<code>memory_internal_free</code>	Memory used as reported by the guest agent (KiB). Enabled by default	None	Free Memory
<code>runnable_any</code>	Fraction of time that at least one vCPU is runnable in the domain	None	vCPUs runnable any
<code>runnable_vcpus</code>	Fraction of time that vCPUs of the domain are runnable	None	vCPUs runnable sum
<code>runstate_fullrun</code>	Fraction of time that all vCPUs are running.	None	vCPUs full run
<code>runstate_fullcontention</code>	Fraction of time that all vCPUs are runnable (that is, waiting for CPU)	None	vCPUs full contention
<code>runstate_concurrencyhazard</code>	Fraction of time that some vCPUs are running and some are runnable	None	vCPUs concurrency hazard
<code>runstate_blocked</code>	Fraction of time that all vCPUs are blocked or offline	None	vCPUs idle

Metric Name	Description	Condition	XenCenter Name
<code>runstate_partial_run</code>	Fraction of time that some vCPUs are running, and some are blocked	None	vCPUs partial run
<code>runstate_partial_contention</code>	Fraction of time that some vCPUs are runnable and some are blocked	None	vCPUs partial contention
<code>vbd_<vbd>_write</code>	Writes to device <code>vbd</code> in bytes per second. Enabled by default	VBD <code>vbd</code> exists	Disk <code>vbd</code> Write
<code>vbd_<vbd>_read</code>	Reads from device <code>vbd</code> in bytes per second. Enabled by default.	VBD <code>vbd</code> exists	Disk <code>vbd</code> Read
<code>vbd_<vbd>_write_latency</code>	Writes to device <code>vbd</code> in microseconds.	VBD <code>vbd</code> exists	Disk <code>vbd</code> Write Latency
<code>vbd_<vbd>_read_latency</code>	Reads from device <code>vbd</code> in microseconds.	VBD <code>vbd</code> exists	Disk <code>vbd</code> Read Latency
<code>vbd <vbd>_iops_read</code>	Read requests per second.	At least one plugged VBD for non-ISO VDI on the host	Disk <code>vbd</code> Read IOPS
<code>vbd <vbd>_iops_write</code>	Write requests per second.	At least one plugged VBD for non-ISO VDI on the host	Disk <code>vbd</code> Write IOPS
<code>vbd <vbd>_iops_total</code>	I/O requests per second.	At least one plugged VBD for non-ISO VDI on the host	Disk <code>vbd</code> Total IOPS
<code>vbd <vbd>_iowait</code>	Percentage of time waiting for I/O.	At least one plugged VBD for non-ISO VDI on the host	Disk <code>vbd</code> IO Wait
<code>vbd <vbd>_inflight</code>	Number of I/O requests currently in flight.	At least one plugged VBD for non-ISO VDI on the host	Disk <code>vbd</code> Inflight Requests
<code>vbd <vbd>_avgqu_sz</code>	Average I/O queue size.	At least one plugged VBD for non-ISO VDI on the host	Disk <code>vbd</code> Queue Size

Metric Name	Description	Condition	XenCenter Name
<code>vif_<vif>_rx</code>	Bytes per second received on virtual interface number <code>vif</code> . Enabled by default.	VIF <code>vif</code> exists	<code>vif</code> Receive
<code>vif_<vif>_tx</code>	Bytes per second transmitted on virtual interface <code>vif</code> . Enabled by default.	VIF <code>vif</code> exists	<code>vif</code> Send
<code>vif_<vif>_rx_errors</code>	Receive errors per second on virtual interface <code>vif</code> . Enabled by default.	VIF <code>vif</code> exists	<code>vif</code> Receive Errors
<code>vif_<vif>_tx_errors</code>	Transmit errors per second on virtual interface <code>vif</code> Enabled by default.	VIF <code>vif</code> exists	<code>vif</code> Send Errors

Note:

The value of `<XenCenter-pif-name>` can be any of the following:

- NIC `<pif>` - if `<pif>` contains `pif_eth#`, where `##` is 0–9
- `<pif>` - if `<pif>` contains `pif_eth#.``##` or `pif_xenbr##` or `pif_bond##`
- `<Internal>` Network `<pif>` - if `<pif>` contains `pif_xapi##`, (note that `<Internal>` appears as is)
- TAP `<tap>` - if `<pif>` contains `pif_tap##`
- xapi Loopback - if `<pif>` contains `pif_lo`

Custom fields and tags

XenCenter supports the creation of tags and custom fields, which allows for organization and quick searching of VMs, storage and so on. For more information, see [Monitoring System Performance](#).

Custom searches

XenCenter supports the creation of customized searches. Searches can be exported and imported, and the results of a search can be displayed in the navigation pane. For more information, see [Monitoring System Performance](#).

Determine throughput of physical bus adapters

For FC, SAS and iSCSI HBAs you can determine the network throughput of your PBDs using the following procedure.

1. List the PBDs on a host.
2. Determine which LUNs are routed over which PBDs.
3. For each PBD and SR, list the VBDs that reference VDIs on the SR.
4. For all active VBDs that are attached to VMs on the host, calculate the combined throughput.

For iSCSI and NFS storage, check your network statistics to determine if there is a throughput bottleneck at the array, or whether the PBD is saturated.

Monitor CPU usage

December 17, 2025

The optimum number of vCPUs per pCPU on a host depends on your use case. During operation, ensure that you monitor the performance of your XenServer® environment and adjust your configuration accordingly.

Terms

In this area, there are various terms that are sometimes used interchangeably. In this article, we use the following terms and meanings:

- **CPU (physical CPU):** The physical hardware attached to a processor socket.
- **Core:** A physical processing unit, capable of one independent thread of execution, which contains all functional units required to support that execution.
- **Hyperthread:** A physical processing unit, capable of one independent thread of execution, which shares some functional units with another hyperthread (also known as its “sibling thread”).
- **Logical CPU (pCPU):** A unit capable of an independent thread of execution that includes a set of registers and an instruction pointer. In a system with hyperthreads enabled, this is a hyperthread. In other cases, it’s a core.
- **Host pCPUs:** The total number of logical CPUs in the host.
- **vCPU (Virtual CPU):** A virtualized logical CPU. This is a logical unit capable of an independent thread of execution, provided to VMs. In XenServer, vCPUs can “time-share” pCPUs, using a scheduler to determine which vCPU is running on which pCPU at any given time.
- **Guest vCPUs:** The vCPUs that are presented to a guest operating system inside a VM.

- **Dom0 vCPUs:** The vCPUs that are visible to the XenServer control domain (dom0).
- **Host total vCPUs:** The sum of dom0 vCPUs and all the guest vCPUs in the host.

General behavior

The total number of vCPUs on a host is the number of vCPUs used by dom0 added to the total number of vCPUs assigned to all the VMs on the host. As you increase the number of vCPUs on a host, you can experience the following types of behavior:

- When the total number of vCPUs on the host is *less than or equal to* the number of pCPUs on the host, the host always provides as much CPU as is requested by the VMs.
- When the total number of vCPUs on the host is *greater than* the number of pCPUs on the host, the host shares the time of the host pCPUs to the VMs. This behavior does not generally affect the VMs because their vCPUs are usually idle for some of the time and, in most cases the host does not reach 100% pCPU usage.
- When the total number of vCPUs on the host is *greater than* the number of pCPUs on the host and the host is *sometimes* reaching 100% host pCPU usage, the vCPUs of the VMs don't receive as much pCPU as they request during the spikes. Instead, during these spikes the VMs slow down to receive a share of the available pCPU on the host.
- When the total number of vCPUs on the host is *greater than* the number of pCPUs on the host and the host is *often* reaching 100% host pCPU usage, the vCPUs of the VMs are continuously slowed down to receive a share of the available CPUs on the host. If the VMs have real-time requirements, this situation is not ideal and you can address it by reducing the number of vCPUs on the host.

The optimum number of vCPUs on a host can depend on the VM users' perception of the speed of their VMs, especially when the VMs have real-time requirements.

Getting information about your CPUs

To find the total number of pCPUs on your host, run the following command:

```
1 xe host-cpu-info --minimal
```

To find the total number of vCPUs (guest and dom0) currently on your host, run the following command:

```
1 xl vcpu-list | grep -v VCPU | wc -l
```

Monitoring CPU usage with RRD metrics

XenServer provides RRD metrics that describe how the vCPUs on your VMs are performing.

When host pCPU usage is 100%

When a host is reaching 100% of host pCPU usage, use these VM metrics to decide whether to move the VM to another host:

runstate_concurrency_hazard

- **runstate_concurrency_hazard > 0%** indicates that sometimes, at least one vCPU is running while at least one other vCPU wants to run but can't get pCPU time. If the vCPUs must coordinate, this behavior causes performance issues.
- **runstate_concurrency_hazard approaching 100%** is a situation to avoid.

Suggested actions:

If there are performance issues, take one of the following actions:

- Decrease the number of vCPUs in the VM.
- Move the VM to another host.
- Decrease the total number of vCPUs on the host by migrating other VMs or decreasing their number of vCPUs.

runstate_partial_contention

- **runstate_partial_contention > 0%** indicates both that at least one vCPU wants to run but can't get pCPU time, and also that at least one other vCPU is blocked (either because there's nothing to do or it's waiting for I/O to complete).
- **runstate_concurrency_hazard approaching 100%** is a situation to avoid.

Suggested action:

Check whether the back end I/O storage servers are overloaded by looking at the back-end metrics provided by your storage vendor. If the storage servers are not overloaded and there are performance issues, take one of the following actions:

- Decrease the number of vCPUs in the VM.
- Move the VM to another host.
- Decrease the total number of vCPUs on the host by migrating other VMs or decreasing their number of vCPUs.

runstate_full_contention

- **runstate_full_contention > 0%** indicates that sometimes the vCPUs want to run all at the same time but none can get pCPU time.
- **runstate_full_contention approaching 100%** is a situation to avoid.

Suggested actions:

If there are performance issues, take one of the following actions:

- Decrease the number of vCPUs in the VM.
- Move the VM to another host.
- Decrease the total number of vCPUs on the host by migrating other VMs or decreasing their number of vCPUs.

When host pCPU usage is less than 100%

If a host is not reaching 100% of host pCPU usage, use these VM metrics to decide whether a VM has the right number of vCPUs:

runstate_fullrun

- **runstate_fullrun = 0%** indicates that the vCPUs are never being used all at the same time.

Suggested action:

Decrease the number of vCPUs in this VM.

- **0% < runstate_fullrun < 100%** indicates that the vCPUs are sometimes being used all at the same time.
- **runstate_fullrun = 100%** indicates that the vCPUs are always being used all at the same time.

Suggested action:

You can increase the number of vCPUs in this VM, until `runstate_fullrun < 100%`. Do not increase the number of vCPUs further, otherwise it can increase the probability of concurrency hazard if the host reaches 100% of pCPU usage.

runstate_partial_run

- **runstate_partial_run = 0%** indicates that either all vCPUs are always being used (full-run=100%) or no vCPUs are being used (idle=100%).
- **0% < runstate_partial_run < 100%** indicates that, sometimes, at least one vCPU is blocked, either because they have nothing to do, or because they are waiting for I/O to complete.

- **runstate_partial_run=100%** indicates that there is always at least one vCPU that is blocked.

Suggested action:

Check whether the back-end I/O storage servers are overloaded. If they are not, the VM probably has too many vCPUs and you can decrease the number of vCPUs in this VM. Having too many vCPUs in a VM can increase the risk of the VM going into the concurrency hazard state when the host CPU usage reaches 100%.

NUMA optimization in XenServer 9

February 19, 2026

Overview

Non-Uniform Memory Access (NUMA) is a multiprocessor memory architecture in which memory is divided into nodes attached to specific CPU sockets. Accessing memory on the local NUMA node has lower latency than accessing memory on a remote node. When VM vCPUs and their memory are co-located on the same NUMA node, workloads can achieve lower memory-access latency and better performance.

UMA and NUMA compared

- UMA (Uniform Memory Access) hosts present a single physical memory node to CPUs.
- NUMA hosts expose multiple memory nodes. CPUs have lower latency when accessing memory from their local node than from remote nodes.

Benefits of optimized NUMA placement

- Improved application responsiveness and workload throughput.
- Higher VM density per host, reducing the need for additional hosts and overall infrastructure cost.

Limitations

- If no single NUMA node on a host has enough free memory for a VM, XS9 starts the VM with memory spread across multiple nodes to allow the VM to start.


- When memory is spread, the VM uses multiple NUMA nodes (`numa_nodes > 1`) and is not NUMA-optimized.
- Performance improvements vary by workload, benchmark, and CPU type.

Monitor and troubleshoot

Per-VM RRD metrics (disabled by default) help monitor NUMA behavior and vCPU scheduling.

Metrics

- *NUMA nodes* (`numa_nodes`): Number of NUMA nodes used by the VM. Value 1 indicates optimized; values greater than 1 indicate not optimized.
- *Memory NUMA node X* (`memory_numa_node_X`): Amount of VM memory allocated from physical NUMA node X. Value 0 means no memory from that node.
- *vCPUs runnable any* (`runnable_any`): Percentage of time at least one vCPU is ready but no pCPU is available. This is an early indicator of CPU contention.
- *vCPUs runnable sum* (`runnable_vcpus`): Expected percentage slowdown across all vCPUs due to scheduling delays.
- *NUMA non-affine vCPUs* (`numa_node_nonaffine_vcpus`): Percentage of time vCPUs run on pCPUs that are not in the same NUMA node as the VM memory. Higher values indicate increased memory-access latency.



Debug checklist

1. Check `NUMA nodes` to see how many NUMA nodes the VM uses.
2. Check `vCPUs runnable any` for domain-level scheduling contention.
3. Check `vCPUs runnable sum` for vCPU-level scheduling delays.
4. Check `NUMA non-affine vCPUs` for NUMA affinity violations and related latency.

Enable metrics and export samples

Per-VM RRD metrics are disabled by default. To enable a metric (for example, `runnable_any`) for a VM and export a sample:

```
1 # Enable the data source for a VM
2 xe vm-data-source-record uuid=<vm-uuid> data-source=runnable_any
3
4 # Export sample RRD output
5 rrd2csv VM:AVERAGE::runnable_any
```

Configure NUMA placement (xe CLI)

NUMA placement optimization is enabled by default on XS9 hosts. No action is required to turn it on.

To disable NUMA placement optimization on a host, set the host NUMA affinity policy to `any`.

```
1 # View current policy
2 xe host-param-get uuid=<host-uuid> param-name=numa-affinity-policy
3
4 # Disable optimization
5 xe host-param-set uuid=<host-uuid> numa-affinity-policy=any
6
7 # Verify
8 xe host-param-get uuid=<host-uuid> param-name=numa-affinity-policy
```

When enabled by default, the policy value is `default_policy`. When disabled, the value is `any`.

Alerts

April 30, 2026

You can configure XenServer® to generate alerts based on any of the available host or VM metrics. In addition, XenServer provides preconfigured alerts that trigger when hosts undergo certain conditions and states.

View alerts

You can view these alerts using XenCenter® or the xe CLI:

- To view alerts in XenCenter, click **Notifications** and then **Alerts**.

The **Alerts** view displays various types of alerts, including performance alerts, system alerts, software update alerts, license alerts, and VM anti-affinity group alerts. For more information, see [Alerts](#).

- To view alerts using the xe CLI, type `xe message-list`.

You can narrow down the list of alerts by specifying additional parameters, for example:

- Specify the priority of the alerts you want to list: `xe message-list priority=<priority>`
- Specify whether it is a pool-related or VM-related alert: `xe message-list class=<Pool|VM>`
- Specify the name of the alert type: `xe message-list name=<alert_name>`. For a list of some of the available alert types, see System alerts.

You can also configure your environment to send alerts as an email. For more information, see [Receive alerts through email](#).

Performance alerts

Performance alerts can be generated when one of the following values exceeds a specified threshold on a managed host, VM, or storage repository (SR): CPU usage, network usage, memory usage, control domain memory usage, storage throughput, or VM disk usage.

By default, the alert repeat interval is set to 60 minutes, it can be modified if necessary. Alerts are displayed on the **Alerts** page in the **Notifications** area in XenCenter. You can also configure XenCenter to send an email for any specified performance alerts along with other serious system alerts.

Any customized alerts that are configured using the xe CLI are also displayed on the **Alerts** page in XenCenter.

Each alert has a corresponding priority/severity level. You can modify these levels and optionally choose to receive an email when the alert is triggered. The default alert priority/severity is set at 3

.

Priority	Name	Description	Default Email Alert
1	Critical	Act now or data may be permanently lost/corrupted.	Yes
2	Major	Act now or some services may fail.	Yes
3	Warning	Act now or a service may suffer.	Yes
4	Minor	Notice that something just improved.	No
5	Information	Day-to-day information (VM Start, Stop, Resume and so on)	No
?	Unknown	Unknown error	No

Configure performance alerts by using XenCenter

1. In the **Resources** pane, select the relevant host, VM, or SR, then click the **General** tab and then **Properties**.
2. Select the **Alerts** tab. The following table summarizes which alerts are available for hosts, VMs, or SRs:

Alert name	Host	VM	SR	Description
Generate CPU usage alerts	X	X		Set the CPU usage and time threshold that trigger the alert.
Generate control domain CPU usage alerts	X			Set the control domain CPU usage and time threshold that trigger the alert.

Alert name	Host	VM	SR	Description
Generate memory usage alerts	X			Set the memory usage and time threshold that trigger the alert.
Generate control domain memory usage alerts	X			Set the control domain memory usage and time threshold that trigger the alert.
Generate control domain free memory alerts	X			Set the control domain free memory and time threshold that trigger the alert.
Generate disk usage alerts		X		Set the disk usage and time threshold trigger the alert.

Alert name	Host	VM	SR	Description
Generate storage throughput alerts			X	Set the storage throughput and time threshold that trigger the alert. Note: Physical Block Devices (PBD) represent the interface between a specific XenServer host and an attached SR. When the total read/write SR throughput activity on a PBD exceeds the threshold you have specified, alerts are generated on the host connected to the PBD. Unlike other XenServer host alerts, this alert must be configured on the SR.
Generate network usage alerts	X	X		Set the network usage and time threshold that trigger the alert.

To change the alert repeat interval, enter the number of minutes in the **Alert repeat interval** box. When an alert threshold has been reached and an alert generated, another alert is not generated until after the alert repeat interval has elapsed.

3. Click **OK** to save your changes.

For comprehensive details on how to view, filter and configure severities for performance alerts, see [Configuring Performance Alerts](#) in the XenCenter documentation.

Configure performance alerts by using the xe CLI

Note:

Triggers for alerts are checked at a minimum interval of five minutes. This interval avoids placing excessive load on the system to check for these conditions and reporting of false positives. Setting an alert repeat interval smaller than five minutes results in the alerts still being generated at the five minute minimum interval.

The performance monitoring `perfmon` tool runs once every five minutes and requests updates from XenServer which are averages over one minute. These defaults can be changed in `/etc/sysconfig/perfmon`.

The `perfmon` tool reads updates every five minutes of performance variables running on the same host. These variables are separated into one group relating to the host itself, and a group for each VM running on that host. For each VM and host, `perfmon` reads the parameter `other-config:perfmon` and uses this string to determine which variables to monitor, and under which circumstances to generate a message.

For example, the following shows an example of configuring a VM “CPU usage” alert by writing an XML string into the parameter `other-config:perfmon`:

```
1 xe vm-param-set uuid=vm_uuid other-config:perfmon=\
2
3 '<config>
4   <variable>
5     <name value="cpu_usage"/>
6     <alarm_trigger_level value="0.5"/>
7   </variable>
8 </config>'
```

Note:

You can use multiple variable nodes.

After setting the new configuration, use the following command to refresh `perfmon` for each host:

```
1 xe host-call-plugin host=host_uuid plugin=perfmon fn=refresh
```

If this refresh is not done, there is a delay before the new configuration takes effect, since by default, `perfmon` checks for new configuration every 30 minutes. This default can be changed in `/etc/sysconfig/perfmon`.

Valid VM elements

- **name**: The name of the variable (no default). If the name value is either `cpu_usage`, `network_usage`, or `disk_usage`, the `rrd_regex` and `alarm_trigger_sense` parameters are not required as defaults for these values are used.
- **alarm_priority**: The priority of the alerts generated (default 3).
- **alarm_trigger_level**: The level of value that triggers an alert (no default).
- **alarm_trigger_sense**: The value is `high` if `alarm_trigger_level` is a maximum value otherwise `low` if the `alarm_trigger_level` is a minimum value (the default `high`).
- **alarm_trigger_period**: The number of seconds that values (above or below the alert threshold) can be received before an alert is sent (the default is 60).
- **alarm_auto_inhibit_period**: The number of seconds this alert will be disabled after an alert is sent (the default is 3600).
- **consolidation_fn**: Combines variables from `rrd_updates` into one value. For `cpu_usage` the default is `average`, for `fs_usage` the default is `get_percent_fs_usage` and for all others - `sum`.
- **rrd_regex**: Matches the names of variables from `xe vm-data-sources-list uuid=vm_uuid`, to compute performance values. This parameter has defaults for the named variables:
 - `cpu_usage`
 - `memory_internal_free`
 - `network_usage`
 - `disk_usage`

If specified, the values of all items returned by `xe vm-data-source-list` whose names match the specified regular expression are consolidated using the method specified as the `consolidation_fn`.

Valid host elements

- **name**: The name of the variable (no default).
- **alarm_priority**: The priority of the alerts generated (default 3).
- **alarm_trigger_level**: The level of value that triggers an alert (no default).
- **alarm_trigger_sense**: The value is `high` when `alarm_trigger_level` is a maximum value otherwise `low` if the `alarm_trigger_level` is a minimum value. (default `high`)
- **alarm_trigger_period**: The number of seconds that values (above or below the alert threshold) can be received before an alert is sent (default 60).

- **alarm_auto_inhibit_period**: The number of seconds that the alert is disabled for after an alert is sent. (default 3600).
- **consolidation_fn**: Combines variables from **rrd_updates** into one value (default **sum** or **average**)
- **rrd_regex**: A regular expression to match the names of variables returned by the **xe vm -data-source-list uuid=vm_uuid** command to use to compute the statistical value. This parameter has defaults for the following named variables:
 - **cpu_usage**
 - **network_usage**
 - **memory_free_kib**
 - **sr_io_throughput_total_xxxxxxx** (where **xxxxxxx** is the first eight characters of the SR-UUID).

SR Throughput: Storage throughput alerts must be configured on the SR rather than the host. For example:

```
1 xe sr-param-set uuid=sr_uuid other-config:perfmon=\
2 '<config>
3   <variable>
4     <name value="sr_io_throughput_total_per_host"/>
5     <alarm_trigger_level value="0.01"/>
6   </variable>
7 </config>'
```

Generic example configuration The following example shows a generic configuration:

```
1 <config>
2   <variable>
3     <name value="NAME_CHOSEN_BY_USER"/>
4     <alarm_trigger_level value="THRESHOLD_LEVEL_FOR_ALERT"/>
5     <alarm_trigger_period value="
6       RAISE_ALERT_AFTER_THIS_MANY_SECONDS_OF_BAD_VALUES"/>
7     <alarm_priority value="PRIORITY_LEVEL"/>
8     <alarm_trigger_sense value="HIGH_OR_LOW"/>
9     <alarm_auto_inhibit_period value="
10      MINIMUM_TIME_BETWEEN_ALERT_FROM_THIS_MONITOR"/>
11     <consolidation_fn value="FUNCTION_FOR_COMBINING_VALUES"/>
12     <rrd_regex value="REGULAR_EXPRESSION_TO_CHOOSE_DATASOURCE_METRIC"/>
13   </variable>
14   ...
15 </variable>
16
17   ...
18 </config>
```

System alerts

The following table displays the system events/conditions that trigger an alert to be displayed on the **Alerts** page in XenCenter.

Name	Priority/Severity	Description
license_expires_soon	2	XenServer License agreement expires soon.
ha-statefile_lost	2	Lost contact with the high availability Storage Repository, act soon.
ha-heartbeat_approaching_timeout	5	High availability approaching timeout, host may reboot unless action is taken.
ha_statefile_approaching_timeout	5	High availability approaching timeout, host may reboot unless action is taken.
haxapi_healthcheck_approaching_timeout	5	High availability approaching timeout, host may reboot unless action is taken.
ha_network_bonding_error	3	Potential service loss. Loss of network that sends high availability heartbeat.
ha_pool_overcommitted	3	Potential service loss. High availability is unable to guarantee protection for configured VMs.
ha_poor_drop_in_plan_exists_for	3	High availability coverage has dropped, more likely to fail, no loss present yet.
ha_protected_vm_restart_failed	2	Service Loss. High availability was unable to restart a protected VM.
ha_host_failed	3	High availability detected that a host failed.
ha_host_was_fenced	4	High availability rebooted a host to protect against VM corruption.

Name	Priority/Severity	Description
redo_log_healthy	4	The XAPI redo log has recovered from a previous error.
redo_log_broken	3	The XAPI redo log has encountered an error.
ip_configured_pif_can_unplug	3	An IP configured NIC can be unplugged by XAPI when using high availability, possibly leading to high availability failure.
host_sync_data_failed	3	Failed to synchronize XenServer performance statistics.
host_clock_skew_detected	3	The host clock is not synchronized with other hosts in the pool.
host_clock_went_backwards	1	The host clock is corrupted.
pool_master_transition	4	A new host has been specified as pool coordinator.
pbd_plug_failed_on_server_start	3	The host failed to connect to Storage at boot time.
auth_external_init_failed	2	The host failed to enable external AD authentication.
auth_external_pool_non-homogeneous	2	Hosts in a pool have different AD authentication configuration.
multipath_period_alert	3	A path to an SR has failed or recovered.
bond-status-changed	3	A link in a bond has disconnected or reconnected.

Software update alerts

- **XenCenter old:** XenServer expects a newer version but can still connect to the current version
- **XenCenter out of date:** XenCenter is too old to connect to XenServer
- **XenServer out of date:** XenServer is an old version that the current XenCenter cannot connect to

- **License expired alert:** XenServer license has expired
- **Missing IQN alert:** XenServer uses iSCSI storage but the host IQN is blank
- **Duplicate IQN alert:** XenServer uses iSCSI storage, and there are duplicate host IQNs

Receive alerts through email

You can configure XenServer to send email notifications when XenServer hosts generate alerts. The mail-alarm utility in XenServer uses sSMTP to send these email notifications. You can enable basic email alerts by using XenCenter or the xe Command Line Interface (CLI).

Use an SMTP server that does not require authentication. Emails sent through SMTP servers that require authentication cannot be delivered.

Enable email alerts by using XenCenter

1. In the **Resources** pane, right-click on a pool and select **Properties**.
2. In the **Properties** window, select **Email Options**.
3. Select the **Send email alert notifications** check box. Enter your preferred destination address for the notification emails and SMTP server details.
4. Choose your preferred language from the **Mail language** list. The default language for performance alert emails is English.

Enable email alerts by using the xe CLI

To configure email alerts, specify your preferred destination address for the notification emails and SMTP server:

```
1 xe pool-param-set uuid=pool_uuid other-config:mail-destination=joe.bloggs@example.com
2 xe pool-param-set uuid=pool_uuid other-config:ssmtp-mailhub=smtp.example.com:<port>
```

XenServer automatically configures the sender address as `noreply@<hostname>`. However, you can set the sender address explicitly:

```
1 xe pool-param-set uuid=pool_uuid other-config:mail-sender=serveralerts@example.com
```

When you turn on email notifications, you receive an email notification when an alert with a priority of 3 or higher is generated. Therefore, the default minimum priority level is 3. You can change this default with the following command:

```
1 xe pool-param-set uuid=pool_uuid other-config:mail-min-priority=level
```

Note:

Some SMTP servers only forward mails with addresses that use FQDNs. If you find that emails are not being forwarded it might be for this reason. In which case, you can set the server host name to the FQDN so this address is used when connecting to your mail server.

To configure the language for the performance alert emails:

```
1 xe pool-param-set uuid=pool_uuid other-config:mail-language=ja-JP
```

The default language for performance alert emails is English.

Monitor host and dom0 resources with NRPE

December 17, 2025

Note:

The NRPE feature is available for XenServer Premium or Trial Edition customers. For more information about XenServer licensing, see [Licensing](#). To upgrade, or to get a XenServer license, visit the [XenServer website](#).

Users with the Pool Admin role can use any third-party monitoring tool that supports the Nagios Remote Plugin Executor (NRPE) to monitor resources consumed by your XenServer® host and dom0 - the control domain of your host.

You can use the following check plugins to monitor host and dom0 resources:

Metric	NRPE check name	Description	Default warning threshold	Default critical threshold	Performance data returned
Host CPU Load	check_host_load	Gets and checks the current load per physical CPU of the host, where load refers to the number of vCPU(s) in a running or runnable state.	3	4	Current system load of the CPU of the host (calculated by taking the average load of the physical CPU of the host).
Host CPU Usage (%)	check_host_cpu	Gets and checks the current average overall CPU usage of the host.	80%	90%	The percentage of host CPU that is currently free and the percentage that is in use.
Host Memory Usage (%)	check_host_memory	Gets and checks the current memory usage of the host.	80%	90%	The percentage of host memory that is currently free and the percentage that is in use.
Host vGPU Usage (%)	check_vgpu	Gets and checks all the current running Nvidia vGPU usage of the host.	80%	90%	The percentage of running vGPU that is currently free and the percentage that is in use.

Metric	NRPE check name	Description	Default warning threshold	Default critical threshold	Performance data returned
Host vGPU Memory Usage (%)	check_vgpu_memory	Gets and checks all the current running Nvidia vGPU memory usage (including the shared memory and graphic memory) of the host.	80%	90%	The percentage of running vGPU memory (including the shared memory and graphic memory) that is currently free and the percentage that is in use.
Dom0 CPU Load	check_load	Gets and checks the current system load average per CPU of dom0, where load refers to the number of processes in a running or runnable state.	2.7,2.6,2.5	3.2,3.1,3	Host CPU load data calculated by taking the average of the last 1, 5, and 15 minutes.
Dom0 CPU Usage (%)	check_cpu	Gets and checks the current average overall CPU usage of dom0.	80%	90%	The average overall CPU usage of dom0 as a percentage.

Metric	NRPE check name	Description	Default warning threshold	Default critical threshold	Performance data returned
Dom0 Memory Usage (%)	check_memory	Gets and checks the current memory usage of dom0.	80%	90%	The percentage of dom0 memory that is currently free and the percentage that is in use.
Dom0 Free Swap (%)	check_swap	Gets and checks the current swap usage of dom0.	20%	10%	The percentage of MB on dom0 that is currently free.
Dom0 Root Partition Free Space (%)	check_disk_root	Gets and checks the current root partition usage of dom0.	20%	10%	The percentage of MB on the dom0 root partition that is currently free.
Dom0 Log Partition Free Space (%)	check_disk_log	Gets and checks the current log partition usage of dom0.	20%	10%	The percentage of MB on the dom0 log partition that is currently free.
Toolstack Status	check_xapi	Gets and checks the status of the XenServer management toolstack (also known as XAPI).			XAPI elapsed uptime in seconds.

Metric	NRPE check name	Description	Default warning threshold	Default critical threshold	Performance data returned
Multipath Status	check_multipath	Gets and checks the status of the storage paths.			The status of the storage paths. OK indicates that all paths are active, WARNING indicates that some paths have failed but more than one path is active, CRITICAL indicates that there is only one path active or that all paths have failed, UNKNOWN indicates that host multipathing is disabled and that the status of the paths cannot be fetched.

NRPE is an on-premises service that runs in dom0 and listens on TCP port (default) 5666 for check execution requests from a monitoring tool. After a request arrives, NRPE parses it, finds the corresponding check command including the parameter's details from the configuration file, and then runs it. The result of the check is sent to the monitoring tool, which stores the results of past checks and provides a graph showing the historical performance data.

Prerequisites

To be able to use NRPE to monitor host and dom0 resources, the monitoring tool you are using must meet the following prerequisites:

- The monitoring tool must be compatible with NRPE version 4.1.0.
- To allow communication between NRPE and the monitoring tool, the monitoring tool must support TLS 1.2 with ciphers [ECDHE-RSA-AES256-GCM-SHA384](#) and [ECDHE-RSA-AES128-GCM-SHA256](#), and the EC curve is [secp384r1](#).

Constraints

- You can configure NRPE settings for an entire pool or for a standalone host that is not part of a pool. Currently, you cannot configure NRPE settings for an individual host in a pool.
- If you add a host to a pool that already has NRPE enabled and configured on it, XenCenter® does not automatically apply the pool's NRPE settings to the new host. You must reconfigure NRPE settings on the pool after adding the new host or configure the new host with same NRPE settings before adding it to the pool.

Note:

When reconfiguring NRPE settings on a pool after adding a new host, ensure the host is up and running.

- If a host is removed from a pool with NRPE enabled and configured on it, XenCenter does not alter the NRPE settings on the host or the pool.

Configure NRPE by using the xe CLI

You can configure NRPE by using the xe CLI or XenCenter. For more information on how to configure NRPE by using XenCenter, see [Monitoring host and dom0 resources with NRPE](#).

After making configuration changes to NRPE, restart the NRPE service by using:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=restart
```

Enable NRPE

NRPE is disabled by default in XenServer. To enable NRPE on a host's control domain (dom0), run the following commands in the xe CLI:

1. Get the host UUID of the host that you want to monitor:

```
xe host-list
```

2. Enable NRPE on the host:

```
xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=enable
```

If the operation runs successfully, this command outputs **Success**. When XenServer restarts, NRPE starts automatically.

To stop, start, restart, or disable NRPE:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=<operation>
```

where **operation** is **stop**, **start**, **restart**, or **disable**.

Monitoring servers

This is a comma-delimited list of IP addresses or host names that are allowed to talk to the NRPE daemon. Network addresses with a bit mask (for example 192.168.1.0/24) are also supported.

View the current list of monitoring servers:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=get-config  
  args:allowed_hosts
```

Allow the monitoring tool to execute checks:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=set-config  
  args:allowed_hosts=<IP address or hostname>
```

Query all NRPE settings:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=get-config
```

Configure multiple NRPE settings:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=set-config  
  args:allowed_hosts=<IP address or hostname> args:ssl_logging=<SSL  
  log level> args:debug=<debug log level>
```

Logs

Debug logging By default, debug logging is disabled.

To check whether debug logging is enabled, run the following command:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=get-config  
  args:debug
```

If **debug**: 0 is returned, debug logging is disabled.

To enable debug logging:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=set-config  
  args:debug=1
```

SSL logging By default, SSL logging is disabled:

```
1 ssl_logging=0x00
```

To check whether SSL logging is enabled, run the following command:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=get-config  
  args:ssl_logging
```

To enable SSL logging:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=set-config  
  args:ssl_logging=0x2f
```

Warning and critical thresholds

For some of these check plugins, you can set warning and critical threshold values so that if the value returned by a check plugin exceeds the threshold values, an alert is generated. The warning threshold indicates a potential issue and the critical threshold indicates a more serious issue that requires immediate attention. Although default values are set for the warning and critical thresholds, you can adjust the threshold values.

To query the default warning and critical threshold values for all the checks, run the following xe CLI command which returns a list of all the checks and their associated warning and critical thresholds:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=get-threshold
```

You can also query the threshold values for a specific check. For example, to get the warning and critical threshold values for the **check_memory** check plugin, run the following xe CLI command:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=get-threshold  
  args:check_memory
```

You can also change the default value of a threshold. For example, to change the default threshold values for the **check_memory** check plugin, run the following xe CLI command:

```
1 xe host-call-plugin host-uuid=<host uuid> plugin=nrpe fn=set-threshold  
  args:check_memory args:w=75 args:c=85
```

Monitor host and dom0 resources with SNMP

March 3, 2026

Note:

The SNMP feature is available for XenServer Premium or Trial Edition customers. For more information about XenServer licensing, see [Licensing](#). To upgrade, or to get a XenServer license, visit the [XenServer website](#).

With the Pool Admin role, you can use SNMP to remotely monitor resources consumed by your XenServer® host and dom0 - the control domain of your host. An SNMP manager, also known as a network management system (NMS), sends query requests to an SNMP agent running on a XenServer host. The SNMP agent replies to these query requests by sending data collected on various metrics back to the NMS. The data that can be collected is defined by object identifiers (OIDs) in a text file called a management information base (MIB). An OID represents a specific piece of measurable information about a network device, such as CPU or memory usage.

You can also configure traps, which are agent-initiated messages that alert the NMS that a specific event has occurred in XenServer. Both query requests and traps can be used to monitor the status of your XenServer pools. These are defined as metric and trap objects and are identified by OIDs in a MIB file `XENSERVER-MIB.txt`, available to download from the [XenServer Downloads page](#). The following tables provide information about these metric and trap objects.

Metric objects

You can request a specific piece of information about your XenServer hosts by using the metrics listed in the following table. These metrics are used by the SNMP manager when sending query requests to an SNMP agent and so you can view this data in your NMS.

You can view the returned data from these metric objects from your NMS or from the xe CLI. To query the metric objects from the xe CLI, run `host-data-source-query` or `vm-data-source-query` and provide the RRDD data source as a value for the `data-source` parameter. For example:

```
1 xe host-data-source-query data-source=cpu_avg host=<host UUID>
```

Note:

By default, the NMS sends OID query requests to SNMP agents using port 161.

Object identifier (OID)	RRDD data source	Returned data	Type
1.3.6.1.4.1.60953.1.1.1.1	<code>memory</code>	Dom0 total memory in MB	Unsigned32
1.3.6.1.4.1.60953.1.1.1.2	<code>memory_internal_free</code>	Dom0 free memory in MB	Unsigned32
1.3.6.1.4.1.60953.1.1.1.3	<code>cpu_usage</code>	Dom0 CPU usage as a percentage	Float
1.3.6.1.4.1.60953.1.1.1.4	<code>memory_total_kib</code>	Host total memory in MB	Unsigned32
1.3.6.1.4.1.60953.1.1.1.5	<code>memory_free_kib</code>	Host free memory in MB	Unsigned32
1.3.6.1.4.1.60953.1.1.1.6	<code>cpu_avg</code>	Host CPU usage as a percentage	Float
1.3.6.1.4.1.60953.1.1.1.7	(see note 1)	pCPUs number	Unsigned32
1.3.6.1.4.1.60953.1.1.1.8	<code>running_vcpus</code>	Running vCPUs number	Unsigned32
1.3.6.1.4.1.60953.1.1.1.9	<code>running_domains</code>	Running VMs number	Unsigned32

Notes:

1. The name of a pCPU is in the format `cpu` followed by a number. To query the number of pCPUs from the xe CLI, run the following command:

```
xe host-data-source-list host=<host UUID> | grep -E 'cpu[0-9]+'
```

This returns a list of the CPU metrics that match the regular expression `cpu[0-9]+`.

Traps

Traps are alerts sent by the SNMP agent to notify the SNMP manager when certain events occur, allowing you to monitor your XenServer hosts and identify issues early. You can configure your SNMP settings to generate a trap when a limit is reached (for example, if the host CPU usage is too high). When a trap is generated, it is sent to your NMS and the following fields are returned as part of the trap object.

Note:

By default, the SNMP agent on the pool coordinator host sends traps to the NMS using UDP port

162.

Object identifier (OID)	Field name	Type	Description
1.3.6.1.4.1.60953.1.10.1.1	operation	String	Can be one of the following values: add or del . operation is add if a trap is generated by XenServer and sent to your NMS (an alert is also created in XenCenter) or del if an alert is destroyed (for example, if you dismiss an alert).
1.3.6.1.4.1.60953.1.10.1.2	ref	String	The reference for the trap object.
1.3.6.1.4.1.60953.1.10.1.3	uuid	String	The UUID of the trap object.
1.3.6.1.4.1.60953.1.10.1.4	name	String	The name of the trap object.
1.3.6.1.4.1.60953.1.10.1.5	priority	Integer	The severity of the trap. Can be one of the following values: 1: Critical, 2: Major, 3: Warning, 4: Minor, 5: Information, others : Unknown.
1.3.6.1.4.1.60953.1.10.1.6	class	String	The category of the trap generated. Can be one of the following values: VM , Host , SR , Pool , VMPP , VMSS , PVS_proxy , VDI , or Certificate .
1.3.6.1.4.1.60953.1.10.1.7	obj-uuid	String	The xapi object UUID of the various classes of the field class .

Object identifier (OID)	Field name	Type	Description
1.3.6.1.4.1.60953.1.10.1.8	<code>timestamp</code>	String	The time at which the trap is generated.
1.3.6.1.4.1.60953.1.10.1.9	<code>body</code>	String	Detailed information about the field <code>name</code> .

Prerequisites

- All hosts in a pool must be running the same XenServer version and this version must contain the SNMP plugin.

Note:

If you cannot see the **SNMP** tab in XenCenter®, it might be because the host or a member of the pool is not running a version of XenServer that supports SNMP. Update the host or pool to the latest version of XenServer.

- The NMS you are using must support SNMPv2c or SNMPv3.
- Your NMS and XenServer must be network-connected.

Constraints

- You can configure SNMP settings for an entire pool or for a standalone host that is not part of a pool. Currently, you cannot configure SNMP settings for an individual host in a pool.
- If you add a host to a pool that already has SNMP enabled and configured on it, XenCenter does not automatically apply the pool's SNMP settings to the new host. You must reconfigure SNMP settings on the pool after adding the new host or configure the new host with the same SNMP settings before adding it to the pool.

Note:

When reconfiguring SNMP settings on a pool after adding a new host, ensure the host is up and running and not in maintenance mode.

- When the SNMP agent is offline, traps cannot be generated. For example, if the SNMP agent is restarted or the pool coordinator is rebooted or re-designated.

Configure SNMP by using the xe CLI

You can configure SNMP by using the xe CLI or XenCenter. For more information on how to configure SNMP by using XenCenter, see [Monitoring host and dom0 resources with SNMP](#).

result objects

When configuring SNMP, all responses are returned in JSON format. If a command executes successfully, it returns the key value pair `"code": 0`. Some commands (such as the `get-config` command) return a nested JSON object called `result`. The `result` JSON object is also required for the `set-config` command which is used to update the SNMP configuration.

The `result` JSON object is made up of the following objects `common`, `agent`, and `nmss`:

`common`

Field	Allowed values	Default value
<code>enabled</code>	<code>no</code> (disable SNMP service) or <code>yes</code> (enable SNMP service)	<code>no</code>
<code>debug_log</code>	<code>no</code> (disable debug logging) or <code>yes</code> (enable debug logging)	<code>no</code>
<code>max_nmss</code>	N/A (This field is read-only and specifies the max number of supported NMSS)	<code>1</code>

`agent`

Field	Allowed values	Default value
<code>v2c</code>	<code>no</code> (disable SNMPv2c) or <code>yes</code> (enable SNMPv2c)	<code>yes</code>
<code>community</code>	COMMON_STRING_TYPE (see note 1)	<code>public</code>
<code>v3</code>	<code>no</code> (disable v3) or <code>yes</code> (enable v3)	<code>no</code>
<code>user_name</code>	COMMON_STRING_TYPE (see note 1)	

Field	Allowed values	Default value
<code>authentication_password</code>	COMMON_STRING_TYPE where length $\geq$ 8 (see note 1)	
<code>authentication_protocol</code>	MD5 or SHA	
<code>privacy_password</code>	COMMON_STRING_TYPE where length $\geq$ 8 (see note 1)	
<code>privacy_protocol</code>	DES or AES	
<code>engine_id</code>	N/A (This field is read-only and is generated when the SNMP agent starts for the first time)	

nmss

Field	Allowed values	Default value
<code>uuid</code>	NMS UUID (You set this when you configure the NMS trap receiver and this value should be consistent across all hosts in a pool)	
<code>address</code>	NMS IPv4 address or host name (FQDN)	
<code>port</code>	1 to 65535	162
<code>v2c</code>	no (disable SNMPv2c), yes (enable SNMPv2c), or support either SNMPv2c or v3.	yes
<code>community</code>	COMMON_STRING_TYPE (see note 1)	public
<code>v3</code>	no (disable v3), yes (enable v3), or support either SNMPv2c or SNMPv3.	no
<code>user_name</code>	COMMON_STRING_TYPE (see note 1)	
<code>authentication_password</code>	COMMON_STRING_TYPE where length $\geq$ 8 (see note 1)	

Field	Allowed values	Default value
<code>authentication_protocol</code>	MD5 or SHA	
<code>privacy_password</code>	COMMON_STRING_TYPE where length >= 8 (see note 1)	
<code>privacy_protocol</code>	DES or AES	

Notes:

- COMMON_STRING_TYPE refers to a string that meets the following requirements:
 - Any combination of letter, number, hyphen (-), period (.), pound (#), at sign (@), equals (=), colon (:) or underscore characters (_).
 - Length between 6 and 32 inclusive.
- Passwords are not stored in plaintext in any configuration file in XenServer. They are converted to a localized key and stored. The `get-config` command shows the password as a hidden constant comprised of asterisks (*).

Configure the SNMP service**Get the status of the SNMP service:**

```
1 xe host-call-plugin host-uuid=<host-uuid> plugin=snmp fn=status
```

Start, stop, or restart the SNMP service:

```
1 xe host-call-plugin host-uuid=<host-uuid> plugin=snmp fn=<operation>
```

where **operation** is `start`, `stop`, or `restart`.

Get the SNMP configuration details:

```
1 xe host-call-plugin host-uuid=<host-uuid> plugin=snmp fn=get-config
```

If successful, this command returns the key value pair `"code": 0` and the `result` JSON object containing the configuration details of the SNMP service. For example:

```
1 "code": 0,
2   "result": {
3
4     "common": {
5
6       "enabled": "no",
7       "debug_log": "no",
```

```

8      "max_nmss": 1
9    }
10   ,
11   "agent": {
12     "v2c": "yes",
13     "v3": "no",
14     "community": "public",
15     "user_name": "",
16     "authentication_password": "",
17     "authentication_protocol": "",
18     "privacy_password": "",
19     "privacy_protocol": "",
20     "engine_id": "<engine_id>"
21   }
22   ,
23   "nmss": []
24 }

```

Copy the `result` JSON object to your preferred text editor and remove all newline (`\n`) characters from the file. Update the fields with your SNMP configuration details. Configure your NMS by referring to your NMS documentation and specifying values for the fields required for the `nmss` object. For more information, refer to the objects listed above.

To configure the SNMP service, run the `set-config` command and provide the edited `result` JSON object as a parameter value to the `args:config` parameter.

Set the SNMP configuration:

```

1 xe host-call-plugin host-uuid=<host-uuid> plugin=snmp fn=set-config
  args:config='<result>'

```

where **result** is the `result` JSON object returned from the `get-config` command that you copied and edited.

Note:

To configure SNMP for an entire pool, you must run the `set-config` command for each host in the pool.

If the configuration changes are successful, the command returns the key value pair `"code": 0`. If the configuration changes are unsuccessful, the `set-config` command returns one of the following key value pairs which indicate that an error has occurred:

- `"code": 1`: Common error string. For example, an unknown exception.
- `"code": 2`: Error string (parameter is missing).
- `"code": 3`: Returns a `message` object as a list where each element is in the format of `[ field_path, key, value, error string]`.

You can also send a test SNMP trap to your NMS to verify that the specified trap receiver information is correct.

Send a test SNMP trap:

```
1 xe host-call-plugin host-uuid=<host-uuid> plugin=snmp fn=send-test-trap
   args:config='{
2   "nmss": [{
3     "uuid": "<uuid>", "address": "<address>", "port": 162, "v2c": "yes", "v3": "no",
       "community": "public", "user_name": "<user_name>",
       "authentication_password": "<authentication_password>",
       "authentication_protocol": "<authentication_protocol>",
       "privacy_password": "<privacy_password>", "privacy_protocol": "<
4     ] }
5   '}
```

This command sends a test trap to your NMS with the `msg_name` of `TEST_TRAP` and the `msg_body` of `This is a test trap from XenServer pool "<pool name>"` to verify the NMS Trap Receiver configuration.

If you do not receive the test trap, check your SNMP configuration again. If unsuccessful, the `send-test-trap` command also returns one of the following key value pairs which indicate that an error has occurred:

- `"code": 1`: Common error string. For example, an unknown exception.
- `"code": 2`: Error string (parameter is missing).
- `"code": 3`: Returns a `message` object as a list where each element is in the format of `[field_path, key, value, error string]`.
- `"code": 4`: Returns a `message` object as a list where each element is in the format of `[nms address, nms port, error string]`.

Manage virtual machines

June 26, 2026

This section provides an overview of how to create Virtual Machines (VMs) using templates. It also explains other preparation methods, including cloning templates and importing previously exported VMs.

What is a virtual machine?

A Virtual Machine (VM) is a software computer that, like a physical computer, runs an operating system and applications. The VM comprises a set of specification and configuration files backed by the

physical resources of a host. Every VM has virtual devices that provide the same functions as physical hardware. VMs can give the benefits of being more portable, more manageable, and more secure. In addition, you can tailor the boot behavior of each VM to your specific requirements. For more information, see [VM Boot Behavior](#).

XenServer® supports guests with any combination of IPv4 or IPv6 configured addresses.

In XenServer VMs can operate in full virtualized mode. Specific processor features are used to ‘trap’ privileged instructions that the virtual machine carries out. This capability enables you to use an unmodified operating system. For network and storage access, emulated devices are presented to the virtual machine. Alternatively, PV drivers can be used for performance and reliability reasons.

Create VMs

Use VM templates

VMs are prepared from templates. A template is a *gold image* that contains all the various configuration settings to create an instance of a specific VM. XenServer ships with a base set of templates, which are *raw* VMs, on which you can install an operating system. Different operating systems require different settings to run at their best. XenServer templates are tuned to maximize operating system performance.

There are two basic methods by which you can create VMs from templates:

- Using a complete pre-configured template.
- Installing an operating system from a CD, ISO image or network repository onto the appropriate provided template.

[Windows VMs](#) describes how to install Windows operating systems onto VMs.

[Linux VMs](#) describes how to install Linux operating systems onto VMs.

Note:

Templates created by older versions of XenServer can be used in newer versions of XenServer. However, templates created in newer versions of XenServer are not compatible with older versions of XenServer.

Other methods of VM creation

In addition to creating VMs from the provided templates, you can use the following methods to create VMs.

Clone an existing VM You can make a copy of an existing VM by *cloning* from a template. Templates are ordinary VMs which are intended to be used as original copies to create instances of VMs from. A VM can be customized and converted into a template. Ensure that you follow the appropriate preparation procedure for the VM. For more information, see [Preparing for Cloning a Windows VM Using Sysprep](#) and [Preparing to Clone a Linux VM](#).

Note:

Templates cannot be used as normal VMs.

XenServer has two mechanisms for cloning VMs:

- A full copy
- Copy-on-Write

The faster Copy-on-Write mode only writes *modified* blocks to disk. Copy-on-Write is designed to save disk space and allow fast clones, but slightly slows down normal disk performance. A template can be fast-cloned multiple times without slowdown.

Note:

If you clone a template into a VM and then convert the clone into a template, disk performance can decrease. The amount of decrease has a linear relationship to the number of times this process has happened. In this event, the `vm-copy` CLI command can be used to perform a full copy of the disks and restore expected levels of disk performance.

Notes for resource pools If you create a template from VM virtual disks on a shared SR, the template cloning operation is forwarded to any host in the pool that can access the shared SRs. However, if you create the template from a VM virtual disk that only has a local SR, the template clone operation is only able to run on the host that can access that SR.

Import an exported VM You can create a VM by *importing* an existing exported VM. Like cloning, exporting and importing a VM is fast way to create more VMs of a certain configuration. Using this method enables you to increase the speed of your deployment. You might, for example, have a special-purpose host configuration that you use many times. After you set up a VM as required, export it and import it later to create another copy of your specially configured VM. You can also use export and import to move a VM to the XenServer host that is in another resource pool.

For details and procedures on importing and exporting VMs, see [Importing and Exporting VMs](#).

XenServer VM Tools

XenServer VM Tools provide high performance I/O services without the overhead of traditional device emulation.

XenServer VM Tools for Windows

XenServer VM Tools for Windows consist of I/O drivers (also known as paravirtualized drivers or PV drivers) and the Management Agent.

The I/O drivers contain storage and network drivers, and low-level management interfaces. These drivers replace the emulated devices and provide high-speed transport between Windows and the XenServer product family software. While installing a Windows operating system, XenServer uses traditional device emulation to present a standard IDE controller and a standard network card to the VM. This emulation allows the Windows installation to use built-in drivers, but with reduced performance due to the overhead inherent in emulating the controller drivers.

The Management Agent, also known as the Guest Agent, is responsible for high-level virtual machine management features and provides a full set of functions to XenCenter.

Install XenServer VM Tools for Windows on each Windows VM for that VM to have a fully supported configuration, and to be able to use the xe CLI or XenCenter. A VM functions without the XenServer VM Tools for Windows, but performance is hampered when the I/O drivers (PV drivers) are not installed. You must install XenServer VM Tools for Windows on Windows VMs to be able to perform the following operations:

- Cleanly shut down, reboot, or suspend a VM
- View VM performance data in XenCenter
- Migrate a running VM (using live migration or storage live migration)
- Create snapshots with memory (checkpoints) or revert to snapshots
- Multipage support (enabled automatically)

For more information, see [Install XenServer VM Tools for Windows](#).

XenServer VM Tools for Linux

XenServer VM Tools for Linux contain a guest agent that provides extra information about the VM to the host.

You must install the XenServer VM Tools for Linux on Linux VMs to be able to perform the following operations:

- View VM performance data in XenCenter
- Adjust the number of vCPUs on a running Linux VM
- Enable dynamic memory control

Note:

You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Red Hat Enterprise Linux 10 (preview), Rocky Linux 8, Rocky Linux 9, CentOS Stream 9, or CentOS Stream 10 VMs as these operating systems do not support memory ballooning with the Xen® hypervisor.

For more information, see [Install XenServer VM Tools for Linux](#).

Find out the virtualization state of a VM

XenCenter reports the virtualization state of a VM on the VM's **General** tab. You can find out whether or not XenServer VM Tools are installed. This tab also displays whether the VM can install and receive updates from Windows Update. The following section lists the messages displayed in XenCenter:

I/O optimized (not optimized): This field displays whether or not the I/O drivers are installed on the VM.

Management Agent installed (not installed): This field displays whether or not the Management Agent is installed on the VM.

Able to (Not able to) receive updates from Windows Update: specifies whether the VM can receive I/O drivers from Windows Update.

Note:

Windows Server Core 2016 does not support using Windows Update to install or update the I/O drivers. Instead use the XenServer VM Tools for Windows installer provided on the [XenServer Downloads page](#).

Install I/O drivers and Management Agent: this message is displayed when the VM does not have the I/O drivers or the Management Agent installed.

Guest UEFI boot and Secure Boot

XenServer enables the following guest operating systems to boot in UEFI mode:

- Windows 10
- Windows 11

- Windows Server 2016
- Windows Server 2019
- Windows Server 2022
- Windows Server 2025
- Red Hat Enterprise Linux 8
- Red Hat Enterprise Linux 9
- Red Hat Enterprise Linux 10 (preview)
- Ubuntu 22.04
- Ubuntu 24.04
- Rocky Linux 8
- Rocky Linux 9
- SUSE Linux Enterprise 15
- Debian Bookworm 12
- Oracle Linux 8
- CentOS 10 Stream supports UEFI boot mode, but Secure Boot is not currently supported due to an upstream issue.

UEFI boot provides a richer interface for the guest operating systems to interact with the hardware, which can significantly reduce VM boot times. If XenServer supports UEFI boot for your guest operating system, we recommend that you choose this boot mode instead of BIOS.

For these operating systems, XenServer also supports Secure Boot. Secure Boot prevents unsigned, incorrectly signed or modified binaries from being run during boot. On a UEFI-enabled VM that enforces Secure Boot, all drivers must be signed. This requirement might limit the range of uses for the VM, but provides the security of blocking unsigned/modified drivers. If you use an unsigned driver, secure boot fails and an alert is shown in XenCenter. Secure Boot also reduces the risk that malware in the guest can manipulate the boot files or run during the boot process.

You must specify the boot mode when creating a VM. It is not possible to change the boot mode of a VM between BIOS and UEFI (or UEFI Secure Boot) after booting the VM for the first time. However, you can change the boot mode between UEFI and UEFI Secure Boot after the VM is used to troubleshoot potential Secure Boot issues. For more information, see [Troubleshooting](#).

Consider the following when enabling UEFI boot on VMs:

- Ensure that a UEFI-enabled Windows VM has at least two vCPUs. UEFI-enabled Linux VMs do not have this restriction.
- You can import or export a UEFI-enabled VM created on XenServer as an OVA, OVF, or an XVA file. Importing a UEFI-enabled VM from OVA or OVF packages created on other hypervisors is not supported.
- To use PVS-Accelerator with UEFI-enabled VMs, ensure that you are using Citrix Provisioning™ 1906 or later.

- For Windows VMs, use the UEFI settings menu to change the screen resolution of the XenCenter console. For detailed instructions, see [Troubleshooting](#).

Note

UEFI-enabled VMs use NVME and E1000 for emulated devices. The emulation information does not display these values until after you install XenServer VM Tools for Windows on the VM. UEFI-enabled VMs also show as only having 2 NICs until after you install XenServer VM Tools for Windows.

Enabling UEFI boot or UEFI Secure Boot

You can use XenCenter or the xe CLI to enable UEFI boot or UEFI Secure Boot for your VM.

For information about creating a UEFI-enabled VM in XenCenter, see [Create a Windows VM by using XenCenter](#) or [Create a Linux VM by using XenCenter](#).

Using the xe CLI to enable UEFI boot or UEFI Secure Boot When you create a VM, run the following command before booting the VM for the first time:

```
1 xe vm-param-set uuid=<UUID> HVM-boot-params:firmware=<MODE>
2 xe vm-param-set uuid=<UUID> platform:device-model=qemu-upstream-uefi
3 xe vm-param-set uuid=<UUID> platform:secureboot=<OPTION>
```

Where, **UUID** is the VM's UUID, **MODE** is either **BIOS** or **uefi**, and **OPTION** is either 'true' or 'false'. If you do not specify the mode, it defaults to **uefi** if that option is supported for your VM operating system. Otherwise, the mode defaults to **BIOS**. If you do not specify the **secureboot** option, it defaults to 'auto'. For UEFI-enabled VMs, the 'auto' behavior is to enable Secure Boot for the VM.

To create a UEFI-enabled VM from a template supplied with XenServer, run the following command:

```
1 UUID=$(xe vm-clone name-label='Windows 10 (64-bit)' new-name-label='Windows 10 (64-bit)(UEFI)')
2 xe template-param-set uuid=<UUID> HVM-boot-params:firmware=<MODE> platform:secureboot=<OPTION>
```

Do not run this command for templates that have something installed on them or templates that you created from a snapshot. The boot mode of these snapshots cannot be changed and, if you attempt to change the boot mode, the VM fails to boot.

When you boot the UEFI-enabled VM the first time you are prompted on the VM console to press any key to start the installation. If you do not start the operating system installation, the VM console switches to the UEFI shell.

To restart the installation process, in the UEFI console, type the following commands.

```
1 EFI:
2 EFI\BOOT\BOOTX64
```

When the installation process restarts, watch the VM console for the installation prompt. When the prompt appears, press any key.

Disabling Secure Boot

You might want to disable Secure Boot on occasion. For example, some types of debugging cannot be enabled on a VM that is in Secure Boot user mode. To disable Secure Boot, you can use the following command:

```
1 xe vm-param-set uuid=<VM_UUID> platform:secureboot=false
```

It can be re-enabled with:

```
1 xe vm-param-set uuid=<VM_UUID> platform:secureboot=true
```

Keys

For Windows VMs:

UEFI-enabled Windows VMs are provisioned with a PK from an ephemeral private key, the Microsoft KEK, the Microsoft Windows Production PCA, and Microsoft third party keys. The VMs are also provided with an up-to-date revocation list from the UEFI forum. This configuration enables Windows VMs to boot with Secure Boot turned on and to receive automatic updates to the keys and revocation list from Microsoft.

For Linux VMs:

To install third-party drivers in a Linux VM that has Secure Boot enabled, you must create a signing key, add it to the VM as a machine owner key (MOK), and use that key to sign the driver. For more information, see [Install third-party drivers on your Secure Boot Linux VM](#).

Troubleshooting your UEFI and UEFI Secure Boot VMs

For information about troubleshooting your UEFI or UEFI Secure Boot VMs, see [Troubleshoot UEFI and Secure Boot problems](#).

vTPM

XenServer enables you to create a virtual Trusted Platform Module (vTPM) and attach it to your VMs. This virtual TPM does not use or pass through any physical TPMs on your hardware. A physical TPM is not required to be able to create a virtual TPM.

All UEFI boot or UEFI Secure Boot VMs allow you to attach a vTPM.

A VM has a one-to-one relationship with its linked vTPM. A VM can only have one vTPM and a vTPM can only be associated with a single VM. Users with the VM Admin and above roles can create and destroy vTPM instances.

Applications running on the VM can access the vTPM through the TPM 2.0 compliant API. TPM 1.2 is not supported. Users with the VM Operator and above roles can access the vTPM through the VM.

To check if your VM has a linked vTPM, in XenCenter go to the **General** tab and look in the **Device Security** section.

Constraints

The following constraints currently apply to VMs created with an attached vTPM:

- vTPM is not supported for VMs in BIOS boot mode.
- While you can export your VMs to OVF/OVA format, any data in your vTPM is lost as part of this process. This lost data can cause the VM to show unexpected behaviors or can prevent it from starting. If you are using any vTPM features in the VM, do not export your VMs using this format.
- BitLocker is not currently supported for VMs with a vTPM attached.
- vTPM operations performed by the user or Windows in the background might fail if the pool coordinator is offline.
- If a VM is running a custom workload that uses the vTPM, and that workload writes a large amount of data or makes a large number of write requests to the vTPM, this might impact the performance of the XenServer toolstack.

Attach a vTPM to a VM

For new VMs, the vTPM can be added during VM creation. For more information, see the documentation for your preferred VM creation method.

If you have an existing UEFI or UEFI Secure Boot VM that you want to add a vTPM to, you can do it by using XenCenter or by using the xe CLI.

By using XenCenter

1. Shut down the VM.
2. Add a vTPM to the VM.
 - a) Right-click on the VM and select **Manage vTPMs**. Or, on the main menu bar, go to **VM > Manage vTPMs**. The **TPM Manager** dialog opens.
 - b) In the **TPM Manager** dialog, add a vTPM.
3. To verify that the VM has a linked vTPM, select the VM and go to its **General** tab and look in the **Device Security** section.
4. Start the VM.

By using the xe CLI

1. Shut down the VM:

```
1 xe vm-shutdown uuid=<vm_uuid>
```

2. Create a vTPM and attach it to the VM:

```
1 xe vtpm-create vm-uuid=<vm_uuid>
```

3. Start the VM:

```
1 xe vm-start uuid=<vm_uuid>
```

Supported guests and allocating resources

For a list of supported guest operating systems, see [Supported Guests, Virtual Memory, and Disk Size Limits](#)

This section describes the differences in virtual device support for the members of the XenServer product family.

XenServer product family virtual device support

The current version of the XenServer product family has some general limitations on virtual devices for VMs. Specific guest operating systems may have lower limits for certain features. The individual guest installation section notes the limitations. For detailed information on configuration limits, see [Configuration Limits](#).

Factors such as hardware and environment can affect the limitations. For information about supported hardware, see the XenServer [Hardware Compatibility List](#).

VM block devices XenServer emulates an IDE bus in the form of an `hda` device. When using Windows, installing the XenServer VM Tools installs a special I/O driver that works in a similar way to Linux, except in a fully virtualized environment.

CPU features

A pool's CPU feature set can change while a VM is running, for example, when a new host is added to an existing pool or when the VM is migrated to a host in another pool. When a pool's CPU feature set changes, the VM continues to use the feature set which was applied when it was started. To update the VM to use the pool's new feature set, you must restart the VM.

Windows VMs

June 22, 2026

Installing Windows VMs on the XenServer® host requires hardware virtualization support (Intel VT or AMD-V).

Note:

Nested virtualization is not supported for Windows VMs hosted on XenServer.

Basic procedure for creating a Windows VM

The process of installing a Windows on to a VM consists of the following steps:

1. Selecting the appropriate Windows template
2. Choosing the appropriate boot mode
3. Installing the Windows operating system
4. Installing the XenServer VM Tools for Windows (*I/O drivers and the Management Agent*)

Warning:

Windows VMs are supported only when the VMs have the XenServer VM Tools for Windows installed.

Windows VM templates

Windows operating systems are installed onto VMs by cloning an appropriate template using either XenCenter® or the `xe` CLI, and then installing the operating system. The templates for individual guests

have predefined platform flags set which define the configuration of the virtual hardware. For example, all Windows VMs are installed with the ACPI Hardware Abstraction Layer (HAL) mode enabled. If you later change one of these VMs to have multiple virtual CPUs, Windows automatically switches the HAL to multi-processor mode.

The available Windows templates are listed in the following table:

Template Name	Supported boot modes	Description
Windows 10 (64-bit)	BIOS, UEFI, UEFI Secure Boot	Used to install Windows 10 (64-bit)
Windows 11	UEFI, UEFI Secure Boot	Used to install Windows 11
Windows Server 2016	BIOS, UEFI, UEFI Secure Boot	Used to install Windows Server 2016 or Windows Server Core 2016
Windows Server 2019	BIOS, UEFI, UEFI Secure Boot	Used to install Windows Server 2019 or Windows Server Core 2019
Windows Server 2022	BIOS, UEFI, UEFI Secure Boot	Used to install Windows Server 2022 or Windows Server Core 2022
Windows Server 2025	BIOS, UEFI, UEFI Secure Boot	Used to install Windows Server 2025 or Windows Server Core 2025

XenServer supports all SKUs (editions) for the listed versions of Windows.

Attach an ISO image library

The Windows operating system can be installed either from an install CD in a physical CD-ROM drive on the XenServer host, or from an ISO image. See [Create ISO images](#) for information on how to make an ISO image from a Windows install CD and make it available for use.

Create a VM by using XenCenter

To create a Windows VM:

1. On the XenCenter toolbar, click the **New VM** button to open the **New VM** wizard.

The **New VM** wizard allows you to configure the new VM, adjusting various parameters for CPU, storage, and networking resources.

2. Select a VM template.

To begin, choose a suitable VM template. Templates contain the setup information to create a VM with a specific guest operating system, and with the optimum storage, CPU, memory, and virtual network configuration. Various different templates are supplied, and you can add custom templates of your own. For more information, see [Template and BIOS options](#).

When selecting a built-in Windows template, you have the option to manage I/O driver updates for XenServer VM Tools via Windows Update by enabling or disabling the **Allow Windows Update to deliver I/O drivers** checkbox.

For custom templates or snapshots, the **New VM** wizard displays the current value of the `has-vendor-device` field. This field must be set to true to allow delivery of PV drivers through Windows Update.

Note:

If you intend to use your Windows VM as a Citrix Provisioning target or with the reset-on-boot flag enabled, you cannot use any of the automated update mechanisms. In such cases, ensure that the **Allow Windows Update to deliver I/O drivers** checkbox is disabled. For more information, see [Settings for Citrix Provisioning targets or reset-on-boot machines](#).

3. (Optional) Configure BIOS strings.

If the OS that you are installing on your VM is compatible only with the original hardware, check the **Copy host BIOS strings to VM** checkbox. For example, you might use this option for an OS installation CD that was packaged with a specific computer. After you start a VM for the first time, the BIOS strings cannot be changed, so make sure they are correct before proceeding.

To copy BIOS strings using the CLI, see [Install VMs from Reseller Option Kit \(BIOS-locked\) Media](#). Advanced users can set user-defined BIOS strings. For more information, see [User-defined BIOS strings](#).

4. Enter a name and an optional description for the new VM.

5. Choose the source of the OS media to install on the new VM.

Installing from a CD/DVD is the simplest option for getting started.

- a) Choose the default installation source option (DVD drive)
- b) Insert the disk into the DVD drive of the XenServer host

XenServer also allows you to pull OS installation media from a range of sources, including a pre-existing ISO library. An ISO image is a file that contains all the information that an optical disc (CD, DVD, and so on) would contain. In this case, an ISO image would contain the same OS data as a Windows installation CD.

To attach a pre-existing ISO library, click **New ISO library** and indicate the location and type of the ISO library. You can then choose the specific operating system ISO media from the list.

6. In the **Installation Media** tab, you can choose a boot mode for the VM. By default, XenCenter selects the most secure boot mode available for the VM operating system version.

Notes:

- The **UEFI Boot** and **UEFI Secure Boot** options appear grayed out if the VM template you have chosen does not support UEFI boot.
- You cannot change the boot mode after you boot the VM for the first time.

For more information, see [Guest UEFI boot and Secure Boot](#).

7. If required, change the option **Create and attach a new vTPM**.

- For VM operating systems that require a vTPM, the option is selected and cannot be unselected.
- For VM operating systems that do not support a vTPM, the option is grayed out and cannot be selected.
- For VM operating systems that support vTPM, but do not require it, choose whether to attach a vTPM to the VM.

For more information, see [vTPM](#).

8. Select a home server for the VM.

A home server is the host which provides the resources for a VM in a pool. When you nominate a home server for a VM, XenServer attempts to start the VM on that host. If this action is not possible, an alternate host within the same pool is selected automatically. To choose a home server, click **Place the VM on this server** and select a host from the list.

Notes:

- In WLB-enabled pools, the nominated home server isn't used for starting, restarting, resuming, or migrating the VM. Instead, Workload Balancing nominates the best host for the VM by analyzing XenServer resource pool metrics and by recommending optimizations.
- If a VM has one or more virtual GPUs assigned to it, the home server nomination doesn't take effect. Instead, the host nomination is based on the virtual GPU placement policy set by the user.
- During rolling pool upgrade, the home server is not considered when migrating the VM. Instead, the VM is migrated back to the host it was on before the upgrade.

If you do not want to nominate a home server, click **Don't assign this VM a home server**. The VM is started on any host with the necessary resources.

Click **Next** to continue.

9. Allocate processor and memory resources for the VM. For a Windows 10 VM (64-bit), the defaults are 2 virtual CPUs and 4 GB of RAM. You can also choose to modify the defaults. Click **Next** to continue.
10. Assign a virtual GPU. The **New VM** wizard prompts you to assign a dedicated GPU or one or more virtual GPUs to the VM. This option enables the VM to use the processing power of the GPU. With this feature, you have better support for high-end 3D professional graphics applications such as CAD/CAM, GIS, and Medical Imaging applications.
11. Allocate and configure storage for the new VM.

Click **Next** to select the default allocation (32 GB) and configuration, or you might want to do the following extra configuration:

- Change the name, description, or size of your virtual disk by clicking **Edit**.
- Add a new virtual disk by selecting **Add**.

12. Configure networking on the new VM.

Click **Next** to select the default NIC and configurations, including an automatically created unique MAC address for each NIC. Alternatively, you might want to do the following extra configuration:

- Change the physical network, MAC address, or Quality of Service (QoS) priority of the virtual disk by clicking **Edit**.
- Add a new virtual NIC by selecting **Add**.

13. (Optional) If this VM is to be used as a template with Citrix Provisioning or with the `reset-on-boot` flag set, ensure **Start the new VM automatically** is not selected. This enables you to do some required configuration before installing Windows.
14. Review settings, and then click **Create Now** to create the VM and return to the **Search** tab.

An icon for your new VM appears under the host in the **Resources** pane.

15. (Optional) If this VM is intended to be used as a template with Citrix Provisioning or has the `reset-on-boot` flag set, configure the VM before installing Windows.

In the host console, type the following command:

```
1 xe vm-param-set uuid=<uuid> has-vendor-device=false
```

The flag `has-vendor-device=false` ensures that Windows Update does not attempt to install or update the I/O drivers included in the XenServer VM Tools. For more information, see [Settings for Citrix Provisioning targets or reset-on-boot machines](#).

16. On the **Resources** pane, select the VM, and then click the **Console** tab to see the VM console.

17. (Optional) If you want to be able to clone the VM, we recommend that you do not run the Windows first-time setup, known as the Out-Of-Box-Experience (OOBE). Instead, when the OOBE starts on the page that asks for region information, press **Ctrl + Shift + F3** to enter audit mode.

You can then use Sysprep to generalize the VM. For more information see [Prepare to clone a Windows VM by using Sysprep](#).

If you do not intend to clone the VM, continue to the following steps in this procedure.

18. Follow the OS installation screen and make your selections.
19. After the OS installation completes and the VM reboots, Install XenServer VM Tools for Windows.

Create a Windows VM by using the CLI

To create a Windows VM from an ISO repository by using the xe CLI:

Note:

For Windows 10 and Windows 11 VMs, the requirement for a vTPM is specified by the template. You are not required to add anything to the xe CLI commands to set up the vTPM.

1. Create a VM from a template:

```
1 xe vm-install new-name-label=<vm_name> template=<template_name>
```

This command returns the UUID of the new VM.

2. (Optional) Change the boot mode of the VM.

```
1 xe vm-param-set uuid=<uuid> HVM-boot-params:firmware=<mode>
2 xe vm-param-set uuid=<uuid> platform:secureboot=<option>
```

The value of `mode` can be either `BIOS` or `uefi` and defaults to `uefi` if that option is supported for your VM operating system. Otherwise, the mode defaults to `BIOS`. The value of `option` can be set to either `true` or `false`. If you do not specify the Secure Boot option, it defaults to `auto`.

For more information, see [Guest UEFI boot and Secure Boot](#).

3. (Optional) If this VM is intended to be used as a template with Citrix Provisioning or has the `reset-on-boot` flag set, configure the VM before installing Windows.

```
1 xe vm-param-set uuid=<uuid> has-vendor-device=false
```

The flag `has-vendor-device=false` ensures that Windows Update does not attempt to install or update the I/O drivers included in the XenServer VM Tools. For more information, see [Settings for Citrix Provisioning targets or reset-on-boot machines](#).

4. Create an ISO storage repository:

```
1 xe-mount-iso-sr <path_to_iso_sr>
```

5. List all of the available ISOs:

```
1 xe cd-list
```

6. Insert the specified ISO into the virtual CD drive of the specified VM:

```
1 xe vm-cd-add vm=<vm_name> cd-name=<iso_name> device=3
```

7. Start the VM and install the operating system:

```
1 xe vm-start vm=<vm_name>
```

At this point, the VM console is visible in XenCenter.

8. On the XenCenter **Resources** pane, select the VM, and then click the **Console** tab to see the VM console.
9. (Optional) If you want to be able to clone the VM, we recommend that you do not run the Windows first-time setup, known as the Out-Of-Box-Experience (OOBE). Instead, when the OOBE starts on the page that asks for region information, press **Ctrl + Shift + F3** to enter audit mode.

You can then use Sysprep to generalize the VM. For more information see [Prepare to clone a Windows VM by using Sysprep](#).

If you do not intend to clone the VM, continue to the following steps in this procedure.
10. Follow the OS installation screen and make your selections.
11. After the OS installation completes and the VM reboots, install the XenServer VM Tools for Windows.

For more information on using the CLI, see [Command line interface](#).

Install XenServer VM Tools for Windows

XenServer VM Tools for Windows provide high performance I/O services without the overhead of traditional device emulation. For more information about the XenServer VM Tools for Windows and advanced usage, see [XenServer VM Tools for Windows](#).

Before you install the XenServer VM Tools for Windows, ensure that your VM is configured to receive the I/O drivers from Windows Update. Windows Update is the recommended way to receive updates to the I/O drivers. However, if Windows Update is not an available option for your VM, you can also receive updates to the I/O drivers through other means. For more information, see [Update the I/O drivers](#).

To install XenServer VM Tools for Windows:

1. We recommend that you snapshot your VM before installing or updating the XenServer VM Tools.
2. Download the XenServer VM Tools for Windows file from the [XenServer Downloads page](#).
3. Verify your download against the provided SHA-256 value.
4. Copy the file to your Windows VM or to a shared drive that the Windows VM can access.
5. Run the `managementagentxXX.msi` file to begin XenServer VM Tools installation.

```
1 Msiexec.exe /package managementagentxXX.msi
```

6. Follow the prompts in the installer.
 - a) Follow the instructions on the wizard to accept the license agreement and choose a destination folder.
 - b) The wizard displays the recommended settings on the **Installation and Updates Settings** page. For information about customizing these settings, see [XenServer VM Tools for Windows](#).
 - c) Click **Next** and then **Install** to begin the XenServer VM Tools for Windows installation process.
7. Restart the VM when prompted to complete the installation process.

Upgrade the Windows operating system in your VM

Upgrades to VMs are typically required when Microsoft releases a new version or revision, or when moving to a newer version of XenServer.

Before you upgrade your Windows VM

1. If you are updating your operating system to a version of Windows that requires a vTPM (such as Windows 11), you must attach a vTPM to your VM before upgrading its operating system. For more information, see [Attach a vTPM to a VM](#).
2. Upgrade XenServer VM Tools for Windows to the latest version on the VM. For more information, see [XenServer VM Tools for Windows](#).
3. Ensure that any older versions of the XenServer VM Tools for Windows drivers are cleaned up. You can do this by running the following commands in the Windows VM:
 - a) Run the XenServer VM Tools for Windows clean up tool: `<tools-install-directory>\uninstall.exe cleanup`

- b) List all the installed drivers by using the command `pnputil.exe -enum-drivers`.

Any drivers with the **Provider Name** XenServer and the **Published Name** of the form `oemxx.inf` must be removed.

- c) To remove these drivers, run the following command for each driver: `pnputil.exe -delete-driver oemXX.inf`

This command removes only older versions of drivers that are not in use. If the driver is in use, the command correctly fails.

Upgrade the Windows operating system

You can upgrade your Windows VMs in one of the following ways:

- Use Windows Update to upgrade the version of the Windows operating system on your Windows VMs. If you use Windows Update to update your XenServer VM Tools, we recommend that you use this method.
- Use the Windows installation ISO for the newer versions. Windows installation disks typically provide an upgrade option if you boot them on a server which has an earlier version of Windows already installed.

In your Windows VM console, follow the upgrade prompts provided by Windows.

Prepare to clone a Windows VM by using Sysprep

The only supported way to clone a Windows VM is by using the Windows utility `sysprep` to prepare the VM.

The `sysprep` utility changes the local computer SID to make it unique to each computer. The `sysprep` binaries are in the `C:\Windows\System32\Sysprep` folder.

For more information about using Sysprep, see [Sysprep \(Generalize\) a Windows installation](#).

To run sysprep on a Windows VM:

Note:

On Windows 10 and 11, the Windows first-time setup or Out-Of-Box-Experience (OOBE), installs applications (such as AppX) that can interfere with the `sysprep` process. Because of this behavior, when creating a clonable VM, we recommend skipping the first-time setup and starting Windows in audit mode instead.

1. Create a Windows VM.
2. Install Windows.

3. (Optional) When the Out-Of-Box-Experience (OOBE) starts on the page that asks for region information, press **Ctrl + Shift + F3**. Windows starts in audit mode. For more information, see [Boot Windows to Audit mode or OOBE](#).

While not required, we recommend that you exit OOBE to avoid creating an unneeded user account on the image and to avoid issues with third-party application compatibility. If you continue with OOBE, some applications or Windows Update that are installed during OOBE might prevent Sysprep from operating correctly.

4. Install the latest version of XenServer VM Tools for Windows.
5. Install any applications and perform any other configuration required.
6. Run `sysprep` to generalize the VM. This utility shuts down the VM when it completes.

Note:

Do not restart the original, generalized VM (the “source” VM) again after the `sysprep` stage. Immediately convert it to a template afterwards to prevent restarts. If the source VM is restarted, `sysprep` must be run on it again before it can be safely used to make more clones.

To clone a generalized Windows VM:

1. Using XenCenter convert the VM into a template.
2. Clone the newly created template into new VMs as required.
3. When the cloned VM starts, it completes the following actions before being available for use:
 - It gets a new SID and name
 - It runs a setup to prompt for configuration values as necessary
 - Finally, it restarts

Customize Windows VMs using Sysprep

XenServer customizes Windows virtual machines using Microsoft Windows Sysprep. During the process, XenServer uploads a user-provided XML file to the VM and runs the Sysprep executable to apply the configuration. The VM completes customization after rebooting. This mechanism enables programmatic configuration of large Windows VM deployments created from a shared image.

This feature requires Windows PV drivers version 9.4.2 or later. XenServer supports it through both the XenServer API and the XE command-line utility, but not through XenCenter.

To initiate customization, the VM Sysprep API call or XE command passes an XML configuration file (typically named `unattend.xml`) to the Sysprep command on the VM. XenServer creates a temporary virtual CD-ROM containing the file and inserts it into the VM. This CD-ROM appears in XenCenter

with a “SYSPREP” prefix and should not be used for other purposes. XenServer deletes the CD-ROM automatically after customization completes.

XenServer does not validate or inspect the XML file and does not report any issues related to its contents.

API instructions

The VM object exposes a `VM.sysprep` method that accepts the customization XML code as a direct parameter along with a timeout value. For more information, see [VM API Documentation](#).

CLI instructions

```
# xe vm-sysprep uuid=176f8e25-3bba-c5b0-7124-cc06a0edc13a filename=unattend.xml
```

The `xe` command accepts the following arguments:

- The UUID of the virtual machine to be customized. The VM must be in a running state.
- The filename of the sysprep file, which will be passed to the Windows sysprep tools.
- An optional timeout value, specified in seconds. If not provided, the default timeout is 180 seconds.

The CLI command returns when the VM reboots or when the timeout expires. Because the customization duration depends on the XML file’s contents, you may need to adjust the timeout value. If the customization finishes after the timeout, the CLI command reports a failure, even though the process may have completed successfully.

Troubleshooting

The Sysprep API and CLI command check several conditions before initiating customization:

- **Feature availability:** XenServer verifies that the Sysprep feature is enabled in `/etc/xapi.conf`.
- **Driver support:** XenServer confirms that the VM’s drivers support the feature. You can inspect driver compatibility from the CLI on the host where the VM is running.

```
1 [root@eu1-dt034 ~]# xenstore-ls -f | grep sysprep
2 /local/domain/11/control/feature-sysprep = "1"
3 /local/domain/14/control/feature-sysprep = "1"
```

The output above shows that the feature is supported on local domains 11 and 14. The driver version used by these domains can be inspected likewise:

```
1 [root@eu1-dt034 ~]# xenstore-ls -f | grep PVAddons
2 /local/domain/11/attr/PVAddons = ""
3 /local/domain/11/attr/PVAddons/BuildVersion = "178"
4 /local/domain/11/attr/PVAddons/Installed = "1"
5 /local/domain/11/attr/PVAddons/MajorVersion = "9"
6 /local/domain/11/attr/PVAddons/MicroVersion = "2"
7 /local/domain/11/attr/PVAddons/MinorVersion = "4"
8 /local/domain/14/attr/PVAddons = ""
9 /local/domain/14/attr/PVAddons/BuildVersion = "178"
10 /local/domain/14/attr/PVAddons/Installed = "1"
11 /local/domain/14/attr/PVAddons/MajorVersion = "9"
12 /local/domain/14/attr/PVAddons/MicroVersion = "2"
13 /local/domain/14/attr/PVAddons/MinorVersion = "4"
```

This can be also observed using the XE CLI:

```
1 # xe vm-param-get param-name=PV-drivers-version uuid=aee4a4d2-75d9-4
   e9f-41d3-c383713a8ac7
2 major: 9; minor: 4; micro: 2; build: 178; xeniface: XenServer 9.1.13.96
   ; xenvif: XenServer 9.1.14.109 ; xenbus: XenServer 9.1.12.120
```

- **Timeout reported by the API or XE command with Sysprep still in progress.** If Sysprep is still running inside the VM when the timeout is reported, increase the timeout from the default of 180 seconds and try again. If Sysprep has stopped or failed (for example, because of a bad answer file), increasing the timeout does not help. Inspect the Windows Sysprep logs inside the VM (%WINDIR%\System32\Sysprep\Panther\) for details.

Windows VM release notes

There are many versions and variations of Windows with different levels of support for the features provided by XenServer. This section lists notes and errata for the known differences.

General Windows issues

- When installing Windows VMs, start off with no more than three virtual disks. After the VM and XenServer VM Tools for Windows have been installed, you can add extra virtual disks. Ensure that the boot device is always one of the initial disks so that the VM can successfully boot without the XenServer VM Tools for Windows.
- When the boot mode for a Windows VM is BIOS boot, Windows formats the primary disk with a Master Boot Record (MBR). MBR limits the maximum addressable storage space of a disk to 2 TiB. To use a disk that is larger than 2 TiB with a Windows VM, do one of the following things:

- If UEFI boot is supported for the version of Windows, ensure that you use UEFI as the boot mode for the Windows VM.
- Create the large disk as the secondary disk for the VM and select GUID Partition Table (GPT) format.
- Multiple vCPUs are exposed as CPU sockets to Windows guests, and are subject to the licensing limitations present in the VM. The number of CPUs present in the guest can be confirmed by checking the Device Manager. The number of CPUs actually being used by Windows can be seen in the Task Manager.
- The disk enumeration order in a Windows guest might differ from the order in which they were initially added. This behavior is because of interaction between the I/O drivers and the Plug-and-Play subsystem in Windows. For example, the first disk might show up as **Disk 1**, the next disk hot plugged as **Disk 0**, a later disk as **Disk 2**, and then upwards in the expected fashion.
- A bug in the VLC player DirectX back-end replaces yellow with blue during video playback when the Windows display properties are set to 24-bit color. VLC using OpenGL as a back-end works correctly, and any other DirectX-based or OpenGL-based video player works too. It is not a problem if the guest is set to use 16-bit color rather than 24.
- The PV Ethernet Adapter reports a speed of 100 Gbps in Windows VMs. This speed is an artificial hardcoded value and is not relevant in a virtual environment because the virtual NIC is connected to a virtual switch. The Windows VM uses the full speed that is available, but the network might not be capable of the full 100 Gbps.
- If you attempt to make an insecure RDP connection to a Windows VM, this action might fail with the following error message: “This could be due to CredSSP encryption oracle remediation.” This error occurs when the Credential Security Support Provider protocol (CredSSP) update is applied to only one of the client and server in the RDP connection. For more information, see <https://support.microsoft.com/en-gb/help/4295591/credssp-encryption-oracle-remediation-error-when-to-rdp-to-azure-vm>.

XenServer® VM Tools for Windows

June 9, 2026

XenServer VM Tools (formerly Citrix VM Tools or XenServer PV Tools) for Windows provide high performance I/O services without the overhead of traditional device emulation. XenServer VM Tools for Windows consist of I/O drivers (also known as paravirtualized drivers or PV drivers) and the Management Agent.

XenServer VM Tools for Windows must be installed on each Windows VM for the VM to have a fully supported configuration. A VM functions without them, but performance is hampered.

The version of the XenServer VM Tools for Windows is updated independently of the version of XenServer. To be in support, ensure that your XenServer VM Tools for Windows are regularly updated to the latest version, both in your VMs and in any templates that you use to create your VMs. For more information about the latest version of the tools, see [What's new](#).

Install XenServer VM Tools

Note:

To install XenServer VM Tools for Windows on a Windows VM, the VM must be running the Microsoft .NET Framework Version 4.0 or later.

The XenServer VM Tools for Windows are installed by default in the `C:\Program Files\XenServer\XenTools` directory on the VM.

To install XenServer VM Tools for Windows:

1. We recommend that you snapshot your VM before installing or updating the XenServer VM Tools.
2. Download the XenServer VM Tools for Windows file from the [XenServer Downloads page](#).
3. Verify your download against the provided SHA-256 value.
4. Copy the file to your Windows VM or to a shared drive that the Windows VM can access.
5. Run the `managementagentx64.msi` file to begin XenServer VM Tools installation.

```
1 Msiexec.exe /package managementagentx64.msi
```

6. Follow the prompts in the installer.
 - Follow the instructions on the wizard to accept the license agreement and choose a destination folder.
 - Customize the settings on the **Installation and Updates Settings** page.

By default, the wizard displays the following recommended setting:

- **Keep XenServer VM Tools for Windows up to date automatically.**

For some use cases, different update settings are recommended. For more information, see [Update XenServer VM Tools](#).

To configure the update settings, you can make the following changes:

- If you do not want to allow the automatic updating of the Management Agent, unselect the checkbox **Keep XenServer VM Tools for Windows up to date automatically**.

- Click **Next** and then **Install** to begin the XenServer VM Tools for Windows installation process.
7. Restart the VM when prompted to complete the installation process.

Customers who install the XenServer VM Tools for Windows or the Management Agent through RDP might not see the restart prompt as it only appears on the Windows console session. To ensure that you restart your VM (if necessary) and to get your VM to an optimized state, specify the force restart option in RDP. The force restart option restarts the VM only if it is required to get the VM to an optimized state.

Warning:

Installing or upgrading the XenServer VM Tools for Windows can cause the friendly name and identifier of some network adapters to change. Any software which is configured to use a particular adapter might have to be reconfigured following XenServer VM Tools for Windows installation or upgrade.

Silent installation

To install the XenServer VM Tools for Windows silently and to prevent the system from rebooting, run the following command:

```
1 Msiexec.exe /package managementagentx64.msi /quiet /norestart
```

A non-interactive, but non-silent installation can be obtained by running:

```
1 Msiexec.exe managementagentx64.msi /passive
```

Notes:

- The `/quiet` parameter applies to the installation dialogs only, but not to the device driver installation. When the `/quiet` parameter is specified, the device driver installation requests permission to reboot if necessary.
 - When `/quiet /norestart` is specified, the system doesn't reboot after the entire tools installation is complete. This behavior is independent of what the user specifies in the reboot dialog.
 - When `/quiet /forcerestart` is specified, the upgrade or installation process can trigger multiple reboots. This behavior is independent of what the user specifies in the reboot dialog.
 - When the device driver installation requests permission to reboot, a tools installation with the `quiet` parameter specified can still be in progress. Use the Task Manager to confirm whether the installer is still running.

To customize the installation settings, use the following parameters with the silent installation commands:

Parameter	Allowed values	Default	Description
ALLOWAUTOUPDATE	YES or NO	YES	Allow automatic Management Agent updates

For example, to do a silent install of the tools that does not allow future automatic Management Agent updates run one of the following commands:

```
1 Msiexec.exe /package managementagentx64.msi ALLOWAUTOUPDATE=NO /quiet /norestart
```

For interactive, silent, and passive installations, following the next system restart there might be several automated reboots before the XenServer VM Tools for Windows are fully installed. This behavior is also the case for installations with the `/norestart` flag specified. However, for installations where the `/norestart` flag is provided, the initial restart might be manually initiated.

Update XenServer VM Tools

Ensure that your XenServer VM Tools for Windows are regularly updated to the latest version, both in your VMs and in any templates that you use to create your VMs. We recommend that you snapshot your VM before updating the XenServer VM Tools. For more information about the latest version of the tools, see [What's new](#).

Important:

Ensure that all requested VM restarts are completed as part of the update. Multiple restarts might be required. If all requested restarts are not completed, the VM might show unexpected behavior.

XenServer provides automatic update mechanisms for each of its components:

- The Management Agent can automatically update itself and the I/O drivers.
- The I/O drivers can also be updated through Windows Update when selecting a VM template in the **New VM** wizard in XenCenter. This approach is not recommended as updates delivered through Windows Update can cause VM reboots. For details on enabling I/O driver updates through Windows Update, see [Overview of VM creation steps](#).

We recommend installing the tools with the MSI installer and using the built-in Management Agent update mechanism to manage updates. When you install the tools with the installer, the Management Agent update mechanism automatically detects VMs provisioned using Machine Creation Services

(MCS) or Citrix Provisioning, allowing you to apply updates on the golden image without impacting provisioned machines.

Alternatively, you can update either or both of these components with the installer.

The update method to choose for each component of the tools can depend on your environment.

Recommended update settings


For most use cases, we recommend installing the XenServer VM Tools for Windows with the MSI installer and using the Management Agent to keep both itself and the I/O drivers up to date. This approach avoids VM reboots that can occur when updates are delivered through Windows Update.

We recommend the following settings:

1. Enable Management Agent updates.
2. Disable I/O driver updates through Windows Update.

During installation:

The default value of the settings are recommended:



Settings for Citrix Provisioning™ targets or reset-on-boot machines

If you intend to use your Windows VM as a Citrix Provisioning target or with the reset-on-boot flag enabled, you cannot use any of the automated update mechanisms. We recommend that you set the following configuration on the master template that you use to create these VMs:


1. During VM creation, ensure that the `has-vendor-device` flag is set to **false**. This prevents Windows Update from automatically updating the I/O drivers for XenServer VM Tools for Windows. You can set this flag by:

- Using the CLI. For more information, see [Create a Windows VM by using the CLI](#).

2. Disable Management Agent updates.

During installation:

Specify this configuration when first installing the XenServer VM Tools:

**Settings for automatic updates by the Management Agent only**


You can configure the Management Agent to update both itself and the I/O drivers.

Choose this approach if your organization requires that you review any updates before applying them to your Windows VMs. If this is the case, you must also redirect the Management Agent to get its updates from an internal server.

1. Disable I/O driver updates through Windows Update.
2. (Optional) Redirect the Management Agent updates.

During installation:

Specify this configuration when first installing the XenServer VM Tools:



Update the Management Agent

XenServer enables you to update the Management Agent automatically on both new and existing Windows VMs. By default, XenServer allows the automatic updating of the Management Agent. You can customize the Management Agent update settings during XenServer VM Tools for Windows installation. The automatic updating of the Management Agent occurs seamlessly, and does not reboot your VM. In scenarios where a VM reboot is required, a message appears on the Console tab of the VM notifying users about the required action.

You can get the Management Agent updates automatically, provided the Windows VM has access to the internet.

Manage Automatic Updates by using the CLI XenServer enables you to use the command line to manage the automatic updating of the I/O drivers and the Management Agent. You can run `msiexec .exe` with the arguments listed in the following table to specify whether the I/O drivers and the Management Agent are automatically updated. For information about installing XenServer VM Tools for Windows by using `msiexec .exe`, see Silent installation.

Note:

For VMs managed using either PVS or MCS, the Management Agent automatically detects non-persistent machines and disables automated updates on them. This allows you to apply tool updates on the golden image without impacting provisioned machines, and prevents update attempts on non-persistent VMs that would be lost on reset.

Argument	Values	Description
ALLOWAUTOUPDATE	YES/NO	Allow/disallow auto updating of the Management Agent

Enable Management Agent updates To enable automatic updating of the Management Agent on a per-VM basis:

1. On the VM, open a command prompt as an administrator.
2. Run the following command:

```
1 reg.exe ADD HKLM\SOFTWARE\XenServer\XenTools\AutoUpdate /t
   REG_DWORD /v DisableAutoUpdate /d 0
```

To enable automatic updating of the Management Agent on a per-pool basis, run the following command in the host console:

```
1 xe pool-param-set uuid=pooluuid guest-agent-config:auto_update_enabled=
   true
```

Disable Management Agent updates To disable automatic updating of the Management Agent on a per-VM basis:

1. On the VM, open a command prompt as an administrator.
2. Run the following command:

```
1 reg.exe ADD HKLM\SOFTWARE\XenServer\XenTools\AutoUpdate /t
   REG_DWORD /v DisableAutoUpdate /d 1
```

To disable automatic updating of the Management Agent on a per-pool basis, run the following command in the host console:

```
1 xe pool-param-set uuid=pooluuid guest-agent-config:auto_update_enabled=
   false
```

Redirect the Management Agent updates XenServer enables customers to redirect Management Agent updates to an internal web server before they are installed. This redirection allows customers to review the updates before they are automatically installed on the VM.

The Management Agent uses an updates file to get information about the available updates. The name of this updates file depends on the version of the Management Agent that you use:

- For Management Agent 9.4.0 and later use <https://pvupdates.vmd.citrix.com>.

We strongly recommend that you upgrade to the latest version of the Management Agent. However, if you use an older version of the Management Agent, we provide the following updates files for use with older versions:

- For Management Agent 9.2.1.35 to 9.3.3.125 <https://pvupdates.vmd.citrix.com/autoupdate.v1.json>.
- For Management Agent 9.0.0.0 to 9.2.0.27 <https://pvupdates.vmd.citrix.com/updates.v9.json>.

Complete the following steps to redirect the Management Agent updates:

1. Download the updates file.
2. Download the Management Agent MSI files referenced in the updates file.
3. Validate the Management Agent MSI file hashes match what's referenced in the updates file
4. Upload the MSI files to an internal web server that your VMs can access.
5. Update the updates file to point to the MSI files on the internal web server.
6. Upload the updates file to the web server.

Automatic updates can also be redirected on a per-VM or a per-pool basis. To redirect updates on a per-VM basis:

1. On the VM, open a command prompt as an administrator.
2. Run the command

```
1 reg.exe ADD HKLM\SOFTWARE\XenServer\XenTools /t REG_SZ /v  
   update_url /d \  
2   url of the update file on the web server
```

To redirect automatic updating of the Management Agent on a per-pool basis, run the following command:

```
1 xe pool-param-set uuid=pooluuid guest-agent-config:auto_update_url=url  
   of the update file on the web server
```

Update the I/O drivers

You can update the I/O drivers through Windows Update or by using the Management Agent. The **Virtualization state** section on a VM's **General** tab in XenCenter indicates if the VM is eligible to receive updates from Windows Update. Alternatively, you can turn off automatic updates and manually manage I/O driver updates.

Each I/O driver (such as xennet, xenvif, xenvbd, xeniface, and xenbus) has its own version. For details on the latest versions, see What's New.

Enable I/O driver updates through Windows Update To get I/O driver updates automatically from Microsoft Windows Update, the following conditions must be met:

- Windows Update is enabled on the VM.
- The VM has internet access or it can connect to a WSUS proxy server.
- You are not running the Core version of Windows Server, as Windows Server Core does not support using Windows Update to install or update the I/O drivers.

You can enable I/O driver updates through Windows Update when selecting a VM template in the **New VM** wizard in XenCenter. For details on enabling I/O driver updates through Windows Update, see [Overview of VM creation steps](#).

If you do not want to receive I/O driver updates from Windows Update, disable Windows Update on your VM, or specify a group policy.

Disable I/O driver updates through Windows Update You can disable I/O driver updates through Windows Update for both new and existing VMs:

- For new VMs: When creating a new VM using the **New VM** wizard in XenCenter, you can disable automatic I/O driver updates by setting a flag. This is particularly important for Citrix Provisioning targets or reset-on-boot machines, where disabling automatic updates ensures stability. For more information, see [Overview of VM creation steps](#) and Settings for Citrix Provisioning targets or reset-on-boot machines. For advanced users who prefer to manage these settings by using the command line, see [Create a Windows VM by using the CLI](#).
- For existing VMs: If the VM is already deployed, you can disable I/O driver updates by configuring a group policy. This ensures that the VM does not receive automatic updates from Windows Update.

Set the Management Agent to update the I/O drivers For more information, see [Install XenServer VM Tools for Windows](#).

Automatic reboots

Ensure that all requested VM restarts are completed as part of the update. Multiple restarts might be required. If all requested restarts are not completed, you might see unexpected behavior.

You can set a registry key that specifies the maximum number of automatic reboots that are performed when you install the drivers through Device Manager or Windows Update. After you have installed the xenbus driver version 9.1.1.8 or later, the XenServer VM Tools for Windows use the guidance provided

by this registry key. If you have an earlier version of xenbus, you might still experience unexpected behavior. If this occurs, reboot the VM from XenCenter.

To use this feature, we recommend that you set the following registry key as soon as possible: `HKLM\System\CurrentControlSet\Services\xenbus_monitor\Parameters\Autoreboot`. The value of the registry key must be a positive integer. We recommend that you set the number of reboots in the registry key to 3.

When this registry key is set, the XenServer VM Tools for Windows perform as many reboots as are needed to complete the updates or the number of reboots specified by the registry key - whichever value is lower.

Before each reboot, Windows can display an alert for 60 seconds that warns of the upcoming reboot. You can dismiss the alert, but this action does not cancel the reboot. Because of this delay between the reboots, wait a few minutes after the initial reboot for the reboot cycle to complete.

Notes:

This setting is required for headless servers with static IP addresses.

This automatic reboot feature only applies to updates to the Windows I/O drivers through Device Manager or Windows Update. If you are using the Management Agent installer to deploy your drivers, the installer disregards this registry key and manages the VM reboots according to its own settings.

If, after waiting for all reboots to complete, you still experience unexpected behavior, initiate a reboot of the VM from XenCenter® to attempt to restore the configuration.

Setting automatic reboots using the command line On your Windows VM, open a command prompt as Administrator and run the following command:

```
1 reg add HKLM\System\CurrentControlSet\Services\xenbus_monitor\Parameters /v Autoreboot /t REG_DWORD /d <number-of-reboots>
```

Alternatively, if your Windows VM is configured to allow remote access, you can run the following command from an Administrator command prompt on another system:

```
1 reg add \\<windows-name>\HKLM\System\CurrentControlSet\Services\xenbus_monitor\Parameters /v Autoreboot /t REG_DWORD /d <number-of-reboots>
```

For more information, see the [reg add Windows command](#).

Setting automatic reboots using the Registry Editor

1. From the Windows Start menu, open the Registry Editor and navigate to the key: `HKLM\System\CurrentControlSet\Services\xenbus_monitor\Parameters`.

2. Create a new subkey of the type `REG_DWORD` with the name: `Autoreboot`.
3. Set the value of `Autoreboot` to an integer, for example, 3.

Other configuration and queries

Find the I/O driver version

To find out the version of the I/O drivers installed on the VM:

1. Navigate to `C:\Windows\System32\drivers`.
2. Locate the driver from the list.
3. Right-click the driver and select **Properties** and then **Details**.

The **File version** field displays the version of the driver installed on the VM.

Find the Management Agent version

To find out the version of the Management Agent installed on the VM:

1. Navigate to `C:\Program Files\XenServer\XenTools`.
2. Right-click `XenGuestAgent` from the list and click **Properties** and then **Details**.

The **File version** field displays the version of the Management Agent installed on the VM.

Behavior changes in version 9.6.0.19

Starting with XenServer VM Tools for Windows 9.6.0.19, device activation behavior has changed to improve reliability when the VM `has-vendor-device` parameter changes.

XenServer VM Tools for Windows now:

- Uses priority-based device activation
- Automatically manages configuration changes
- Uses consistent child device identifiers across device transitions
- Automatically transitions to the correct PV device after restart

These improvements are particularly important for MCS and PVS environments that use shared golden images.

No manual device cleanup or additional reboot cycles are required after changing the `has-vendor-device` parameter.

A shutdown and restart is still required for the parameter change to take effect.

Troubleshoot XenServer VM Tools for Windows

We provide a tool, [xtbugtool](#), that captures diagnostic information for the XenServer VM Tools. For more information, see [Troubleshoot VM problems](#).

Switching the `has-vendor-device` parameter

Changing the VMs `has-vendor-device` parameter requires the VM to shutdown and restart.

Note:

Switching `has-vendor-device` is not recommended once the image has been deployed as `reset-on-boot` images rebuilds configuration on every boot and this leads to changes in the network configurations.

In XenServer VM Tools for Windows 9.6.0.19 and later:

- The correct PV device is automatically selected during boot
- Duplicate backend connections are prevented
- Device configuration mismatches are automatically resolved

No additional reboot cycles or manual intervention inside the VM should be required after restart.

If older versions of XenServer VM Tools for Windows are installed, additional reboots or manual device cleanup might still be required.

Uninstall XenServer VM Tools

We don't recommend removing the XenServer VM Tools from your Windows VMs. These tools are required for your Windows VMs to be fully supported. Removing them can cause unexpected behavior. Manually uninstall your XenServer VM Tools only as a last resort.

Standard uninstall

To do a standard uninstall of the XenServer VM Tools, you can use the Windows **Add or Remove Programs** feature:

1. Create a snapshot of the VM before you start.
2. In the Windows VM, go to **Add or Remove Programs**.
3. Select **XenServer VM Tools** and click **Uninstall**.
4. Reboot the VM.

The `uninstall.exe` command

Uninstalling the XenServer VM Tools by using the Windows **Add or Remove Programs** feature calls the `<tools-install-directory>\uninstall.exe` file to perform the uninstall actions. You can instead choose to call this command from a PowerShell terminal or a command prompt with administrator privileges.

1. Create a snapshot of the VM before you start.
2. As an administrator, open a command prompt or PowerShell terminal.
3. Run the command `<tools-install-directory>\uninstall.exe`.
4. Reboot the VM.

Command options The `uninstall.exe` command accepts the following parameters:

- `help` - Displays usage information for the command.
- `log` - Generates a log file that indicates what the command has done.
- `verbose` - Prints to the console what the command has done.
- `disable` - Disables drivers that were installed by the installer MSI.
- `force-disable` - Disables the drivers in all situations.
- `hidden` - Deletes hidden devices. These devices are unused and have been superseded, but might have left stale registry entries.
- `cleanup` - Removes old uninstallers from **Add or Remove Programs**. These uninstallers can include duplicate entries from older versions of the tools.
- `purge` - (9.3.1 and later) Resets the VM to a clean state as it was before any part of the XenServer VM Tools were installed. For more information, see Full uninstall of all XenServer VM Tools components.
- `install` - (9.3.1 and later) Installs the current set of I/O drivers and prompts for a VM reboot when required.
- `reboot` - Reboots the VM after all other command operations have completed.

Full uninstall of all XenServer VM Tools components

The latest version of XenServer VM Tools for Windows (9.3.1 and later) includes the command `uninstall.exe purge`. The `purge` option on the `uninstall.exe` application resets a VM to the state before any of the I/O drivers were installed. If you are experiencing issues when upgrading your tools to a newer version or need a clean state to install a later set of tools on your VM, use this utility.

1. Create a snapshot of the VM before you start.
2. As an administrator, open a command prompt or PowerShell terminal.

3. Run the command `<tools-install-directory>\uninstall.exe purge verbose`
4. Reboot the VM.

After using this command, you do not need to perform any manual cleanup steps like you might have had to with previous versions of the XenServer VM Tools. All changes related to the XenServer VM Tools have been removed.

What does the purge option remove? If you use the command `uninstall.exe purge`, all traces of the XenServer VM Tools are removed from your Windows VM. The list of actions taken by this command are as follows:

- Services:
 - Disables all XenServer VM Tools services, which prevents installed drivers and services from starting on reboot.
 - Stops any running XenServer VM Tools services.
- Drivers:
 - Uninstalls I/O drivers from all device nodes.
 - Uninstalls hidden devices. This action is the same as that performed the `hidden` command line option.
 - Uninstalls cached driver packages, which removes them from the driver store. As a result, the I/O drivers are not automatically reinstalled.
- Registry:
 - Removes stale registry information used by out of support versions of the drivers.
 - Deletes tools-related keys from `HKLM\System\CurrentControlSet\Control\Class\...`
 - Deletes tools-related keys from `HKLM\System\CurrentControlSet\Services`.
 - Deletes tools-related keys from `HKLM\System\CurrentControlSet\Enum\...`
- Files:
 - Deletes any XenServer VM Tools driver files from `C:\Windows\System32` and `C:\Windows\System32\drivers`.
 - Deletes XenServer VM Tools `INF` files from `C:\Windows\INF`.
 - Deletes any stale files left by out of support versions of the tools from `C:\Program Files\Citrix\XenTools` and `C:\Program Files\XenServer\XenTools`.
- Other:
 - Deletes old entires in **Add or Remove Programs**. This action is the same as that performed by the `cleanup` command line option.

- Clears some of the InstallAgent's stale state information.
- Removes `xenfilt.sys` from upper-filters. This change prevents `xenfilt.sys` from loading on any driver nodes.
- Removes the `unplug` keys, which causes the VM to revert to emulated devices on reboot.
- Removes StorNvme's StartOverride. This change forces `stornvme.sys` to start on boot and allows emulated NVMe (UEFI) boot devices to function.

Best Practices

Recommended version for MCS and PVS deployments

Use XenServer VM Tools for Windows 9.6.0.19 or later for environments where VMs transition between `has-vendor-device=true` and `has-vendor-device=false`.

This version improves PV device transition reliability and reduces configuration mismatches for shared golden image deployments.

What's new

The version of the XenServer VM Tools for Windows is updated independently of the version of XenServer. Ensure that your XenServer VM Tools for Windows are regularly updated to the latest version, both in your VMs and in any templates that you use to create your VMs.

The latest version of the XenServer VM Tools for Windows is available from the [XenServer Downloads page](#).

XenServer VM Tools for Windows 9.6.0

Released May 13, 2026

This set of tools contains the following component versions:

- Installer: 9.6.0
- Management Agent 9.6.0.12
- xenbus 9.1.13.23
- xeniface 9.1.14.6
- xennet 9.1.8.6
- xenvbd 9.1.11.5
- xenvif 9.1.15.8

Changes in 9.6.0

- Improved resiliency to VM configuration changes.
- Improved notifications from the system tray icon.

Fixed in 9.6.0

- Fixed an issue where RSS network traffic was not getting distributed across CPUs.
- Fixed a bugcheck when updating XenServer VM Tools for Windows, in some configurations.

Earlier releases

9.4.2 Released Sep 24, 2025

This set of tools contains the following component versions:

- Installer: 9.4.2
- Management Agent 9.3.2.95
- xenbus 9.1.12.120
- xeniface 9.1.13.96
- xennet 9.1.7.65
- xenvbd 9.1.10.84
- xenvif 9.1.14.109

Changes in 9.4.2

- Simplified installation options
- System tray icon, providing XenServer VM Tools for Windows notifications

This release also includes updates to collect telemetry.

Fixed in 9.4.2

- Improved distribution of network workload across CPUs

9.4.1 Released May 27, 2025

This set of tools contains the following component versions:

- Installer: 9.4.1
- Management Agent: 9.3.1.80
- xenbus: 9.1.11.115
- xeniface: 9.1.12.94
- xennet: 9.1.7.65
- xenvbd: 9.1.9.82
- xenvif: 9.1.13.107

All customers who are affected by the issues described in [CTX692748 - XenServer and Citrix Hypervisor Security Bulletin](#) should update to this version of the tools or later.

This release also includes updates to collected telemetry.

Fixed issues in 9.4.1 This release contains fixes for the following issues:

- In some cases, with XenServer VM Tools 9.4.0 installed, Windows 11 cannot be upgraded from 23H2 to 24H2.
- The Management Agent can sometimes crash on VM reboot or startup.
- Some Windows VMs can experience unexpected hard reboots.

9.4.0 Released Oct 22, 2024

This set of tools contains the following component versions:

- Installer: 9.4.0
- Management Agent: 9.3.0.65
- xenbus: 9.1.9.105
- xeniface: 9.1.10.87
- xennet: 9.1.7.65
- xenvbd: 9.1.8.79
- xenvif: 9.1.12.101

This release contains fixes for the following issues:

- The XenServer VM Tools for Windows version is not correctly reported to Citrix Hypervisor™ due to the Management Agent checking the incorrect InstallStatus value in the registry.
- Due to an incomplete handling of StorNvme's StartOverride registry keys, you might encounter a BSOD with the error "Inaccessible Boot Device" after rebooting your Windows VM during a XenServer VM Tools update through Windows Update.
- Sometimes, when updating XenServer VM Tools for Windows, the static IP settings are lost and the network settings change to use DHCP. This issue is resolved if you're updating from version 9.3.3 or later (for example, from 9.4.0 to a newer version). However, if you're updating from a version earlier than 9.3.3 (for example, from 9.3.2 to 9.4.0), the issue will occur once during the update, but it will not affect future updates. For more information, see the support article [Static IP loss when updating to VM Tools for Windows 9.3.3 or earlier](#).

This release contains the following known issues:

- If you rely on the Management Agent for automatic updates, you must manually update to version 9.4.0 or later, as 9.3.3 does not automatically update to 9.4.0 due to certificate checking issues. To resolve this, manually install XenServer VM Tools for Windows 9.4.0 or a newer version on each VM.

9.3.3 Released Jun 04, 2024

This set of tools contains the following component versions:

- Installer: 9.3.3
- Management Agent: 9.2.3.60
- xenbus: 9.1.8.88
- xeniface: 9.1.9.78
- xennet: 9.1.6.60
- xenvbd: 9.1.7.66
- xenvif: 9.1.11.92

This release includes fixes for the following issues:

- Sometimes, the Management Agent can crash with a `ManagementException` when the VM is under heavy load.
- On Windows 2022 Servers, a single vCPU can consume 100% CPU. As a result, the server becomes unresponsive.
- Windows VMs hosted in a pool that uses OpenVSwitch networking and VLAN tagged ethernet frames can have inconsistent network connectivity.

9.3.2 Released Nov 27, 2023

This release rebrands the Citrix® VM Tools to XenServer VM Tools.

This set of tools contains the following component versions:

- Installer: 9.3.2
- Management Agent: 9.2.2.435
- xenbus: 9.1.7.80
- xeniface: 9.1.8.69
- xennet: 9.1.5.51
- xenvbd: 9.1.6.58
- xenvif: 9.1.10.83

This release also contains the following improvements:

- Improvements to the `uninstall.exe` command.
- Changes to enable some Windows VMs to use up to 64 vCPUs where your version of XenServer and the Windows operating system support it.

This release includes a fix for the following issue:

- Sometimes, when installing the XenServer VM Tools, a non-fatal error can cause the installation to fail.

9.3.1 Released Jan 25, 2023

This set of tools contains the following component versions:

- Installer: 9.3.1
- Management Agent: 9.2.1.35
- xenbus: 9.1.5.54
- xeniface: 9.1.5.42
- xennet: 9.1.3.34
- xenvbd: 9.1.4.37
- xenvif: 9.1.8.58

This release includes the following improvements:

- Improvements to the `uninstall.exe` utility, including the `purge` parameter. For more information, see Uninstall XenServer VM Tools.
- General improvements to the XenServer VM Tools installer.
- General improvements to string handling of registry keys.

This release contains fixes for the following issues:

- Sometimes, when the XenServer VM Tools are updated through Windows Update, the static IP settings are lost and the network settings change to use DHCP.
- On Windows VMs, the grant tables can easily become exhausted. When this occurs, read and write requests can fail or additional VIFs are not enabled correctly and fail to start.
- On rare occasions, when upgrading the XenServer VM Tools for Windows, the existing Management Agent can fail to shut down and prevent the upgrade from succeeding.
- On a Windows VM, you might see both the previous and a later version of the tools or Management Agent listed in your Installed Programs.
 - (PREVIOUS) Citrix XenServer Windows Management Agent
 - (LATER) Citrix Hypervisor PV Tools.

After you update to the latest version of the tools, neither of these earlier names are listed. Only XenServer VM Tools is listed in your Installed Programs.

9.3.0 Released Jul 26, 2022

This set of tools contains the following component versions:

- Installer: 9.3.0
- Management Agent: 9.2.0.27
- xenbus: 9.1.4.49

- xeniface: 9.1.4.34
- xennet: 9.1.3.34
- xenvbd: 9.1.3.33
- xenvif: 9.1.6.52

This release includes the following improvements:

- General improvements to the XenServer VM Tools installer.

This release contains fixes for the following issues:

- Security software was blocking secondary disks that are marked as removable from being exposed to the OS, as a data-exfiltration prevention measure. This update enables you to flag a VBD as non-removable and have this correctly exposed through the OS.
- On a Windows VM, sometimes the IP address of an SR-IOV VIF is not visible in XenCenter.

9.2.3 Released Apr 28, 2022

This set of tools contains the following component versions:

- Installer: 9.2.3
- Management Agent: 9.1.1.13
- xenbus: 9.1.3.30
- xeniface: 9.1.4.34
- xennet:
 - 9.1.1.8 (for Windows Server 2012 and Windows Server 2012 R2)
 - 9.1.2.23 (for all other supported Windows operating systems)
- xenvbd: 9.1.2.20
- xenvif: 9.1.5.48

This release contains fixes for the following issues:

- In XenServer VM Tools for Windows version 9.2.2, time synchronize options are not available.
- A race condition can cause Windows VMs to show a blue screen error after live migration on Citrix Hypervisor 8.2 Cumulative Update 1.
- Windows VMs that have version 9.2.1 or 9.2.2 of the XenServer VM Tools installed and that are PVS targets can sometimes freeze with a black screen. The message “Guest Rx stalled” is present in the dom0 kernel logs. This issue more often occurs on pool coordinators than on other pool members.
- On Windows VMs with more than 8 vCPUs, Receive Side Scaling might not work because the xenvif driver fails to set up the indirection table.

9.2.2 Released Jan 14, 2022

This set of tools contains the following component versions:

- Installer: 9.2.2
- Management Agent: 9.1.1.13
- xenbus: 9.1.3.30
- xeniface: 9.1.2.22
- xennet:
 - 9.1.1.8 (for Windows Server 2012 and Windows Server 2012 R2)
 - 9.1.2.23 (for all other supported Windows operating systems)
- xenvbd: 9.1.2.20
- xenvif: 9.1.3.31

This release contains fixes for the following issues:

- During an update of the tools, the xenbus driver can prompt a reboot before driver installation is complete. Accepting the reboot can cause a blue screen error in your Windows VM.
- When compressing collected diagnostic information, the xt-bugtool diagnostics tool times out after 20s. This behavior can result in the diagnostics zip file not being correctly created.
- VNC clipboard sharing doesn't work.
- The previous versions of the drivers were not released through Windows Update.

9.2.1 Released Jun 24, 2021

This set of tools contains the following component versions:

- Installer: 9.2.1
- Management Agent: 9.1.0.10
- xenbus: 9.1.2.14
- xeniface: 9.1.1.11
- xennet: 9.1.1.8
- xenvbd: 9.1.1.8
- xenvif: 9.1.2.16

Note:

This set of drivers was not provided through Windows Update.

This release contains fixes for the following issues:

- In some cases, the Laptop/Slate state of the VM cannot be changed.
- After a VM is restarted it can sometimes begin to send excessive log messages to the daemon.log file.

- A race condition in driver load dependencies after an OS upgrade can prevent the XenServer VM Tools from being upgraded.
- A storage error can cause Windows VMs to crash.
- Sometimes the IP address of an SR-IOV VIF is not visible in XenCenter. To fix the issue, restart the Management Agent from within the VM's Service Manager.
- Under high network and system load, and low resources, VMs can experience bugchecks in both Citrix and third party drivers, typically with the code `IRQL_NOT_LESS_OR_EQUAL`. This fix improves network buffering to prevent these bugchecks.
- Upgrading the Windows I/O drivers can cause UEFI VMs to fail to boot, reporting "0xC000000E. A required device isn't connected or can't be accessed.'"
- An issue can occur when installing the XenServer VM Tools after uninstalling a previous version of the XenServer VM Tools that returns the following error message: "This Device cannot start (code 10) (Operation failed) The requested operation was unsuccessful".

Linux VMs

June 26, 2026

When you want to create a Linux VM, create the VM using a template for the operating system you want to run on the VM. You can use a template that XenServer provides for your operating system, or one that you created previously. You can create the VM from either XenCenter® or the CLI. This section focuses on using the CLI.

Note:

To create a VM of a newer minor update of a RHEL release than is supported for installation by XenServer®, complete the following steps:

- Install from the latest supported media
- Use `yum update` to bring the VM up-to-date

This process also applies to RHEL derivatives such as CentOS and Oracle Linux.

We recommend that you install the XenServer VM Tools for Linux immediately after installing the operating system. For more information, see [Install XenServer VM Tools for Linux](#).

The overview for creating a Linux VM is as following:

1. Create the VM for your target operating system using XenCenter or the CLI.
2. Install the operating system using vendor installation media.
3. Install the XenServer VM Tools for Linux (recommended).

4. Configure the correct time and time zone on the VM and VNC as you would in a normal non-virtual environment.

XenServer supports the installation of many Linux distributions as VMs.

Warning:

The **Other install media** template is for advanced users who want to attempt to install VMs running unsupported operating systems. XenServer has been tested running only the supported distributions and specific versions covered by the standard supplied templates. Any VMs installed using the **Other install media** template are *not* supported.

For information regarding specific Linux distributions, see [Installation notes for Linux distributions](#).

Supported Linux distributions

For a list of supported Linux distributions, see [Guest operating system support](#).

Other Linux distributions are **not** supported. However, distributions that use the same installation mechanism as Red Hat Enterprise Linux (for example, Fedora Core) might be successfully installed using the same template.

Create a Linux VM

This section includes procedures for creating a Linux VM by installing the OS from a physical CD/DVD or from a network-accessible ISO.

Create a Linux VM by using the xe CLI

This section shows the CLI procedure for creating a Linux VM by installing the OS from a physical CD/DVD or from a network-accessible ISO.

1. Create a VM from the appropriate template. The UUID of the VM is returned:

```
1 xe vm-install template=template-name new-name-label=vm-name
```

2. (Optional) Change the boot mode of the VM.

```
1 xe vm-param-set uuid=<uuid> HVM-boot-params:firmware=<mode>
2 xe vm-param-set uuid=<UUID> platform:device-model=qemu-upstream-uefi
3 xe vm-param-set uuid=<uuid> platform:secureboot=<option>
```

The value of `mode` can be either `BIOS` or `uefi` and defaults to `uefi` if that option is supported for your VM operating system. Otherwise, the mode defaults to `BIOS`. The value of `option` can be set to either `true` or `false`. If you do not specify the Secure Boot option, it defaults to `auto`.

For more information, see [Guest UEFI boot and Secure Boot](#).

3. Add a virtual CD-ROM to the new VM:

- If you are installing from a CD or DVD, get the name of the physical CD drive on the XenServer host:

```
1 xe cd-list
```

The result of this command gives you something like SCSI 0:0:0:0 for the `name-label` field.

Use this value parameter as the `cd-name` parameter:

```
1 xe vm-cd-add vm=vm_name cd-name="host_cd_drive_name_label"
  device=3
```

- If you are installing from a network-accessible ISO, use the name of the ISO from the ISO library-label as the value for the `cd-name` parameter:

```
1 xe vm-cd-add vm=vm_name cd-name="iso_name.iso" device=3
```

4. Insert the operating system installation CD into the CD drive on the XenServer host.
5. Open a console to the VM with XenCenter or an SSH terminal and follow the steps to perform the OS installation.
6. Start the VM. It boots straight into the operating system installer:

```
1 xe vm-start uuid=UUID
```

7. Install the guest utilities and configure graphical display. For more information, see [Install the XenServer VM Tools for Linux](#).

Create a Linux VM by using XenCenter

1. On the XenCenter toolbar, click the **New VM** button to open the **New VM** wizard.

The **New VM** wizard allows you to configure the new VM, adjusting various parameters for CPU, storage, and networking resources.

2. Select a VM template and click **Next**.

Each template contains the setup information that is required to create a VM with a specific guest operating system (OS), and with optimum storage. This list reflects the templates that XenServer currently supports.

Note:

If the OS that you are installing on your VM is compatible only with the original hardware, check the **Copy host BIOS strings to VM** box. For example, you might use this option for an OS installation CD that was packaged with a specific computer.

After you first start a VM, you cannot change its BIOS strings. Ensure that the BIOS strings are correct before starting the VM for the first time.

To copy BIOS strings using the CLI, see [Install VMs from Reseller Option Kit \(BIOS-locked\) Media](#).

Advanced users can set user-defined BIOS strings. For more information, see [User-defined BIOS strings](#).

3. Enter a name and an optional description for the new VM.
4. Choose the source of the OS media to install on the new VM.

Installing from a CD/DVD is the simplest option for getting started.

- a) Choose the default installation source option (DVD drive)
- b) Insert the disk into the DVD drive of the XenServer host

XenServer also allows you to pull OS installation media from a range of sources, including a pre-existing ISO library.

To attach a pre-existing ISO library, click **New ISO library** and indicate the location and type of the ISO library. You can then choose the specific operating system ISO media from the list.

5. In the **Installation Media** tab, you can choose a boot mode for the VM. By default, XenCenter selects the most secure boot mode available for the VM operating system version.

Notes:

- The **UEFI Boot** and **UEFI Secure Boot** options appear grayed out if the VM template you have chosen does not support UEFI boot.
- You cannot change the boot mode after you boot the VM for the first time.

For more information, see [Guest UEFI boot and Secure Boot](#).

6. Select a home server for the VM.

A home server is the host which provides the resources for a VM in a pool. When you nominate a home server for a VM, XenServer attempts to start the VM on that host. If this action is not possible, an alternate host within the same pool is selected automatically. To choose a home server, click **Place the VM on this server** and select a host from the list.

Notes:

- In WLB-enabled pools, the nominated home server isn't used for starting, restarting, resuming, or migrating the VM. Instead, Workload Balancing nominates the best host for the VM by analyzing XenServer resource pool metrics and by recommending optimizations.
 - If a VM has one or more virtual GPUs assigned to it, the home server nomination doesn't take effect. Instead, the host nomination is based on the virtual GPU placement policy set by the user.
 - During rolling pool upgrade, the home server is not considered when migrating the VM. Instead, the VM is migrated back to the host it was on before the upgrade.
- If you do not want to nominate a home server, click **Don't assign this VM a home server**. The VM is started on any host with the necessary resources.

Click **Next** to continue.

7. Allocate processor and memory resources for the VM. Click **Next** to continue.

8. Assign a virtual GPU.

If vGPU is supported, the **New VM** wizard prompts you to assign a dedicated GPU or one or more virtual GPUs to the VM. This option enables the VM to use the processing power of the GPU. With this feature, you have better support for high-end 3D professional graphics applications such as CAD/CAM, GIS, and Medical Imaging applications.

9. Allocate and configure storage for the new VM.

Click **Next** to select the default allocation (24 GB) and configuration, or you might want to do the following extra configuration:

- Change the name, description, or size of your virtual disk by clicking **Edit**.
- Add a new virtual disk by selecting **Add**.

10. Configure networking on the new VM.

Click **Next** to select the default NIC and configurations, including an automatically created unique MAC address for each NIC. Alternatively, you might want to do the following extra configuration:

- Change the physical network, MAC address, or Quality of Service (QoS) priority of the virtual disk by clicking **Edit**.
- Add a new virtual NIC by selecting **Add**.

11. Review settings, and then click **Create Now** to create the VM and return to the **Search** tab.

An icon for your new VM appears under the host in the **Resources** pane.

On the **Resources** pane, select the VM, and then click the **Console** tab to see the VM console.

12. Follow the OS installation screens and make your selections.
13. After the OS installation completes and the VM reboots, install the XenServer VM Tools for Linux.

Create a Linux VM by using PXE boot

You can use PXE boot to install the operating system of your Linux VM. This approach can be useful when you have to create many Linux VMs.

To install by using PXE boot, set up the following prerequisites in the network where your Linux VMs are located:

- DHCP server that is configured to direct any PXE boot installation requests to the TFTP server
- TFTP server that hosts the installation files for the Linux operating system

When creating the Linux VM, run the following commands:

1. Create a VM from the appropriate template. The UUID of the VM is returned:

```
1 xe vm-install template=template-name new-name-label=vm-name
```

2. Set the boot order to boot from the disk and then from the network:

```
1 xe vm-param-set uuid=<UUID> HVM-boot-params:order=cn
```

3. Start the VM to begin the PXE boot installation:

```
1 xe vm-start uuid=<UUID>
```

4. Install the guest utilities and configure graphical display. For more information, see [Install the XenServer VM Tools for Linux](#).

For more information about using PXE boot to install Linux operating systems, see the operating system documentation:

- Debian: [Installing Debian using network booting](#)
- Red Hat: [Starting a Kickstart installation automatically using PXE](#)
- CentOS: [Product documentation](#)
- SLES: [Preparing Network Boot Environment](#)
- Ubuntu: [Netbooting the server installer on amd64](#)

Install XenServer VM Tools for Linux

Although all supported Linux distributions are natively paravirtualized (and don't need special drivers for full performance), XenServer VM Tools for Linux provide a guest agent. This guest agent provides

extra information about the VM to the host. Install the guest agent on each Linux VM to benefit from the following features:

- View VM performance data in XenCenter.

For example, the following memory performance values are visible in XenCenter only when the XenServer VM Tools are installed: “Used Memory”, “Disks”, Network”and “Address”.

- In XenCenter, view the Linux guest operating system information.
- In the XenCenter **Networking** tab, view the IP address of the VM.
- Launch an SSH Console to the VM from XenCenter.
- Adjust the number of vCPUs on a running Linux VM.
- Enable Dynamic Memory Control (DMC).

Note:

You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Red Hat Enterprise Linux 10 (preview), Rocky Linux 8, Rocky Linux 9, CentOS Stream 9, or CentOS Stream 10 VMs as these operating systems do not support memory ballooning with the Xen® hypervisor.

It is important to keep the Linux guest agent up-to-date as you upgrade your XenServer host. For more information, see [Update Linux kernels and guest utilities](#).

Note:

Before installing the guest agent on a SUSE Linux Enterprise Desktop or Server 15 guest, ensure that `insserv-compat-0.1-2.15.noarch.rpm` is installed on the guest.

To install the XenServer VM Tools for Linux:

1. Download the XenServer VM Tools for Linux file from the [XenServer Downloads page](#).
2. Copy the `LinuxGuestTools-xxx.tar.gz` file to your Linux VM or to a shared drive that the Linux VM can access.
3. Extract the contents of the tar file: `tar -xzf LinuxGuestTools-xxx.tar.gz`
4. Run the installation script as the root user:

```
1 /<extract-directory>/install.sh
```

5. If the kernel has been upgraded, or the VM was upgraded from a previous version, reboot the VM now.

Uninstall XenServer VM Tools for Linux

From version 8.4.0-1, you can use the `install.sh` script to uninstall XenServer VM Tools for Linux. To uninstall the tools, run the following command as the root user:

```
1 /<extract-directory>/install.sh -u
```

Install third-party drivers on your Secure Boot Linux VM

To install third-party drivers in a Linux VM that has UEFI Secure Boot enabled, you must create a signing key, add it to the VM as a machine owner key (MOK), and use that key to sign the driver. For example, if you use the XenServer graphics capabilities with your Linux VM, you might need to install the NVIDIA graphics driver on your VM.

Complete the following steps to create a key and use it to install a third-party driver:

1. Generate a public key and private key pair.
2. Enroll the public key in MOK.
3. Set the keys you created as the module signing keys for the driver.

The following example shows this procedure in detail for an NVIDIA graphics driver on a secure-boot-enabled Ubuntu VM:

1. Download the NVIDIA driver to your VM.
2. Create a directory (for example, `/root/module-signing`) to hold the keys:

```
1 mkdir -p /root/module-signing
```

3. Create a public and private key to use to sign the driver:

```
1 openssl req -new -x509 -newkey rsa:2048 -keyout /root/module-  
signing/Nvidia.key -outform DER -out /root/module-signing/  
Nvidia.der -nodes -days 36500 -subj "/CN=Graphics Drivers"
```

4. Import the public key into MOK by using `mokutil`:

```
1 mokutil --import /root/module-signing/Nvidia.der
```


You are asked to create a password during this step. When you next boot, you are prompted to provide the password you create here.

5. Ensure that the VM boot target is set to graphical:

```
1 systemctl set-default graphical.target
```

6. Reboot the VM.

7. During the boot, the **Perform MOK management** GUI is displayed.



In this interface, complete the following steps:

- Select **Enroll MOK > Continue**.
 - When asked whether to **Enroll the key(s)?**, select **Yes**.
 - When prompted, provide the password you created when you imported the public key (Step 4).
8. Install the package `libglvnd-dev`:

```
1 apt install pkg-config libglvnd-dev
```

9. Install the NVIDIA driver, specifying the keys you created as the module signing keys:

```
1 bash ./NVIDIA-Linux-x86_64-535.129.03-grid.run --module-signing-  
secret-key=/root/module-signing/Nvidia.key --module-signing-  
public-key=/root/module-signing/Nvidia.der
```

Installation notes for Linux distributions

This section lists vendor-specific, configuration information to consider before creating the specified Linux VMs.

For more detailed release notes on all distributions, see [Linux VM Release Notes](#).

Apt repositories (Debian)

For infrequent or one-off installations, it is reasonable to use a Debian mirror directly. However, if you intend to do several VM installations, we recommend that you use a caching proxy or local mirror. Either of the following tools can be installed into a VM.

- **Apt-cacher**: An implementation of proxy server that keeps a local cache of packages
- **debmirror**: A tool that creates a partial or full mirror of a Debian repository

Prepare to clone a Linux VM

Typically, when cloning a VM or a computer, unless you generalize the cloned image, attributes unique to that machine are duplicated in your environments. Some of the unique attributes that are duplicated when cloning are the IP address, SID, or MAC address.

As a result, XenServer automatically changes some virtual hardware parameters when you clone a Linux VM. When you copy the VM using XenCenter, XenCenter automatically changes the MAC address and IP address for you. If these interfaces are configured dynamically in your environment, you might not need to modify the cloned VM. However, if the interfaces are statically configured, you might need to modify their network configurations.

The VM may need to be customized to be made aware of these changes. For instructions for specific supported Linux distributions, see [Linux VM Release Notes](#).

Machine name

A cloned VM is another computer, and like any new computer in a network, it must have a unique name within the network domain.

IP address

A cloned VM must have a unique IP address within the network domain it is part of. Generally, this requirement is not a problem when DHCP is used to assign addresses. When the VM boots, the DHCP server assigns it an IP address. If the cloned VM had a static IP address, the clone must be given an unused IP address before being booted.

MAC address

There are two situations when we recommend disabling MAC address rules before cloning:

1. In some Linux distributions, the MAC address for the virtual network interface of a cloned VM is recorded in the network configuration files. However, when you clone a VM, XenCenter assigns the new cloned VM a different MAC address. As a result, when the new VM is started for the first time, the network does not recognize the new VM and does not come up automatically.
2. Some Linux distributions use udev rules to remember the MAC address of each network interface, and persist a name for that interface. This behavior is intended so that the same physical NIC always maps to the same `ethn` interface, which is useful with removable NICs (like laptops). However, this behavior is problematic in the context of VMs.

For example, consider the behavior in the following case:

```
1 1. Configure two virtual NICs when installing a VM
2 1. Shut down the VM
3 1. Remove the first NIC
```

When the VM reboots, XenCenter shows just one NIC, but calls it `eth0`. Meanwhile the VM is deliberately forcing this NIC to be `eth1`. The result is that networking does not work.

For VMs that use persistent names, disable these rules before cloning. If you do not want to turn off persistent names, you must reconfigure networking inside the VM (in the usual way). However, the information shown in XenCenter does not match the addresses actually in your network.

Update Linux kernels and guest utilities

The Linux guest utilities can be updated by rerunning the `install.sh` script from the XenServer VM Tools for Linux (see [Install the XenServer VM Tools for Linux](#)).

For `yum`-enabled distributions, CentOS and RHEL, `xe-guest-utilities` installs a `yum` configuration file to enable subsequent updates to be done using `yum` in the standard manner.

For Debian, `/etc/apt/sources.list` is populated to enable updates using `apt` by default.

When upgrading, we recommend that you always rerun `install.sh`. This script automatically determines if your VM needs any updates and installs if necessary.

Multipage support for Linux VMs

Multipage support is available for better storage performance.

To enable multipage support on a Linux VM complete the following steps:

1. First, verify that your kernel supports the `max_ring_page_order` parameter.

Run the following command: `modinfo xen_blkfront | grep max_ring_page_order`

If this command returns an empty response, your kernel does not support this feature. Do not proceed with these steps.

2. Take a snapshot of your VM.
3. On the VM, run the following command: `echo 'options xen_blkfront max_ring_page_order=3'>/etc/modprobe.d/xen_blkfront.conf`
4. Depending on your Linux distribution, run one of the following commands:
 - For RHEL, CentOS, or Oracle Enterprise Linux: `dracut -f -v`
 - For Debian-based distributions: `update-initramfs -u -k all`
5. Reboot your VM.

Linux VM release notes

Most modern Linux distributions support Xen paravirtualization directly, but have different installation mechanisms and some kernel limitations.

RHEL graphical install support

To use the graphical installer, in XenCenter step through the **New VM** wizard. In the **Installation Media** page, in the **Advanced OS boot parameters** section, add `vnc` to the list parameters:

```
1 graphical utf8 vnc
```

You are prompted to provide networking configuration for the new VM to enable VNC communication. Work through the remainder of the **New VM** wizard. When the wizard completes, in the **Infrastructure** view, select the VM, and click **Console** to view a console session of the VM. At this point, it uses the standard installer. The VM installation initially starts in text mode, and may request network configuration. Once provided, the **Switch to Graphical Console** button is displayed in the top right corner of the XenCenter window.

Red Hat Enterprise Linux 8

You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Red Hat Enterprise Linux 10 (preview), Rocky Linux 8, Rocky Linux 9, CentOS Stream 9, or CentOS Stream 10 VMs as these operating systems do not support memory ballooning with the Xen hypervisor.

SUSE Linux Enterprise 12

Prepare a SLES guest for cloning

Note:

Before you prepare a SLES guest for cloning, ensure that you clear the udev configuration for network devices as follows:

```
1 cat< /dev/null > /etc/udev/rules.d/30-net_persistent_names.rules
```

To prepare an SLES guest for cloning:

1. Open the file `/etc/sysconfig/network/config`
2. Edit the line that reads:

```
1 FORCE_PERSISTENT_NAMES=yes
```

To

```
1 FORCE_PERSISTENT_NAMES=no
```

3. Save the changes and reboot the VM.
For more information, see [Prepare to Clone a Linux VM](#).

Ubuntu 24.04

- The Hardware Enablement (HWE) kernel on Ubuntu 22.04 is not supported.
- Due to a known issue with Mesa drivers on Ubuntu 24.04, after installing an Ubuntu 24.04 (desktop) VM on an AMD platform host with the Wayland protocol enabled by default, you might experience repeated re-login prompts.

To resolve this issue, you can either disable Wayland or switch to lightdm as the default display manager.

- Disable Wayland.

Open a command prompt and run the following command:

```
1 sudo nano /etc/gdm3/custom.conf
```

Set the following value to disable Wayland:

```
1 WaylandEnable=false
```

Restart the VM to complete the process.

- Switch to lightdm.

Install lightdm:

```
1 sudo apt install lightdm
```

Set it as the default display manager:

```
1 sudo dpkg-reconfigure lightdm
```

Restart the VM to complete the process.

Migrate VMs

May 29, 2026

You can migrate a running VM by using live migration or storage live migration to move a VM's Virtual Disk Image (VDI) without any VM downtime.

Live migration and storage live migration

The following sections describe the compatibility requirements and limitations of live migration and storage live migration.

Live migration

Live migration is available in all versions of XenServer. This feature enables you to move a running VM from one host to another host, when the VM's disks (VDIs) are on storage shared by both hosts. Pool maintenance features such as high availability and Rolling Pool Upgrade (RPU) can automatically move VMs by using live migration. These features allow for workload leveling, infrastructure resilience, and the upgrade of server software, without any VM downtime.

During the live migration of a VM, its memory is transferred as a data stream between two hosts using the network. The migration stream compression feature compresses this data stream, speeding up the memory transfer on slow networks. This feature is disabled by default, but this can be changed by using XenCenter or the xe CLI. For more information, see [Pool Properties - Advanced](#) and [Pool parameters](#). Alternatively, you can enable compression when migrating a VM by using the command line. For more information, see the `vm-migrate` command in [VM Commands](#).

The parallel host evacuation feature speeds up host evacuation time (during host updates) by moving VMs off a host in parallel instead of sequentially. By default, this feature is enabled and the VMs are

migrated in batches of 10 in parallel. You can change the default batch size in the `/etc/xapi.conf` file.

Note:

Storage can only be shared between hosts in the same pool. As a result VMs can only be migrated to hosts in the same pool.

Storage live migration**Notes:**

- Storage migration cannot be used with non-persistent VMs.
- Storage live migration cannot be used on VMs that have changed block tracking enabled. Disable changed block tracking before attempting storage live migration.
- Storage migration (live or offline) cannot be used on VMs whose VDIs are on a GFS2 or XFS SR. However, you can manually copy the VDIs to another SR.

Storage live migration allows a VM to be moved from one host to another, when the VM's disks are not on storage shared between the two hosts. As a result, VMs stored on local storage can be migrated without downtime and VMs can be moved from one pool to another. This feature enables system administrators to:

- Rebalance VMs between XenServer pools (for example from a development environment to a production environment).
- Upgrade and update standalone XenServer hosts without any VM downtime.
- Upgrade XenServer server hardware.

Note:

- Migrating a VM from one host to another preserves the VM *state*. The state information includes information that defines and identifies the VM and the historical performance metrics, such as CPU and network usage.

Compatibility requirements

When migrating a VM with live migration or storage live migration, the VM and servers must meet the following compatibility requirements.

General requirements:

- The target host must have the same or a more recent version of XenServer installed as the source host. For hosts on the same version of XenServer, this requirement also includes ensuring that your target host has the same or more recent level of updates as the source host.

- XenServer VM Tools for Windows must be installed on each Windows VM that you want to migrate.
- You cannot concurrently migrate more than three VMs that have their source location in the same pool.

CPU requirements:

- If the CPUs on the source and target host are different, the target host must be at least as capable as the source host. Generally, this means that the target has the same or a newer CPU.
 - If you are migrating within the same pool, the pool automatically attempts to make a VM compatible.
 - If you are migrating between pools, you must ensure that the VM is compatible with the feature set in the destination pool.
- You cannot live migrate a VM between AMD and Intel processors.

Memory requirements:

- The target host must have sufficient spare memory capacity or be able to free sufficient capacity using Dynamic Memory Control. If there is not enough memory, the migration fails to complete.
- Storage migration only: A host in the source pool must have sufficient spare memory capacity to run a halted VM that is being migrated. This requirement enables the halted VM to be started at any point during the migration process.

Disk space requirements:

- Storage migration only: The target storage must have enough free disk space available for the incoming VMs. The free space required can be three times the VDI size (without snapshots). If there is not enough space, the migration fails to complete.
- The source storage must have enough free disk space to create temporary snapshots of the VM's VDIs during the migration. If there is not enough space, the migration fails to complete. The free space required can be up to two times the size of the VM's disk.

Limitations and caveats

Live migration and storage live migration are subject to the following limitations and caveats. Moving a shut-down VM to another SR by using the **Move VM** wizard is subject to the same limitations as storage migration (live or offline):

- Storage migration cannot be used with VMs created by Machine Creation Services.
- VMs using PCI pass-through devices cannot be migrated (except in the case of NVIDIA SR-IOV GPUs). For more information, see [Use SR-IOV enabled NICs](#).

- VMs with attached vUSBs cannot be migrated.
- VMs with the parameter `no-migrate` set cannot be migrated.
- You cannot use storage live migration to migrate VMs that have changed block tracking enabled. Disable changed block tracking before attempting storage live migration. For more information, see [Changed Block Tracking](#).
- VMs that have the `on-boot` option set to `reset` cannot be migrated. For more information, see [Intellicache](#).
- If you use the high availability feature and the VM being migrated is marked as protected, you might receive a warning during live migration if the operation causes the HA constraints to not be met.
- VM performance is reduced during migration.
- Time to completion of VM migration depends on the memory footprint of the VM, and its activity. In addition, the size of the VDI and the storage activity of the VDI can affect VMs being migrated with storage live migration. VMs with vGPUs attached migrate the entire vGPU state while the VM is paused. We recommend that you use a fast network card on the management network to reduce downtime, especially with vGPUs that have large amounts of memory.
- If live migration fails, for example, in the case of a network error, the VM on the source host can instantly go to a halted state.

Migrate a VM using XenCenter®

1. In the Resources pane, select the VM and do one of the following:
 - To migrate a running or suspended VM using live migration or storage live migration, on the **VM** menu, click **Migrate to Server** and then **Migrate VM wizard....** This action opens the **Migrate VM** wizard.
 - To move a stopped VM: On the **VM** menu, select **Move VM....** This action opens the **Move VM** wizard.
2. In the **Destination Pool** pane, select a standalone server or a pool from the **Destination** list. If you want to place the VM's disks on a server's local storage, select the server from the bottom list. Click **Next**.
3. In the **Storage** pane, specify the storage repository where you would like to place the migrated VM's virtual disks, and then click **Next**.
 - The **Place all migrated virtual disks on the same SR** radio button is selected by default and displays the default shared SR on the destination pool.
 - Click **Place migrated virtual disks onto specified SRs** to specify an SR from the **Storage Repository** list. This option allows you to select different SR for each virtual disk on the migrated VM.

4. If you are migrating the VM across pools, map the VM's virtual network interfaces to networks in the destination pool or standalone server in the **Networking** pane, and then click **Next**.
5. In the **Storage Migration Network** pane, select a network on the destination pool that will be used for the live migration of the VM's virtual disks. Click **Next**.

Note:

Due to performance reasons, it is recommended that you do not use your management network for live migration.

6. If you are migrating the VM across servers within the same pool, use the **Memory Migration Network** pane to select a network that will be used for the live migration of the VM's memory . Click **Next**.

Note:

Due to performance reasons, it is recommended that you do not use your management network for live migration.

7. Review the configuration settings and click **Finish** to start migrating the VM.

If you are upgrading your version of XenServer, you might need to shut down and boot all VMs after migrating your VMs, to ensure that new virtualization features are picked up.

Live VDI migration

Live VDI migration allows the administrator to relocate the VMs Virtual Disk Image (VDI) without shutting down the VM. This feature enables administrative operations such as:

- Moving a VM from cheap local storage to fast, resilient, array-backed storage.
- Moving a VM from a development to production environment.
- Moving between tiers of storage when a VM is limited by storage capacity.
- Performing storage array upgrades.

Limitations and caveats

Live VDI Migration is subject to the following limitations and caveats

- Storage migration cannot be used with non-persistent VMs.
- IPv6 Linux VMs require a Linux Kernel greater than 3.0.
- If you perform live VDI migration on a VM that has a vGPU, vGPU live migration is used. The host must have enough vGPU space to make a copy of the vGPU instance on the host. If the pGPUs are fully employed, VDI migration might not be possible.

- When you do a VDI live migration for a VM that remains on the same host, that VM temporarily requires twice the amount of RAM.

To move virtual disks

1. In the **Resources** pane, select the SR where the Virtual Disk is stored and then click the **Storage** tab.
2. In the **Virtual Disks** list, select the Virtual Disk that you would like to move, and then click **Move**.
3. In the **Move Virtual Disk** dialog box, select the target SR that you would like to move the VDI to.

Note:

Ensure that the SR has sufficient space for another virtual disk: the available space is shown in the list of available SRs.

4. Click **Move** to move the virtual disk.

Import and export VMs

May 29, 2026

XenServer® allows you to import VMs from and export them to various different formats.

Using the XenCenter **Import** wizard, you can import VMs from Open Virtualization Format (OVF and OVA), disk image formats (VHD, VHDX/AVHDX, and VMDK), and XenServer XVA format. You can even import VMs that have been created on other virtualization platforms, such as those offered by VMware and Microsoft.

Note:

When importing VMs that have been created using other virtualization platforms, configure or *fix up* the guest operating system to ensure that it boots on XenServer. The Operating System Fixup feature in XenCenter aims to provide this basic level of interoperability. For more information, see [Operating system fixup](#).

Using the XenCenter® **Export** wizard, you can export VMs to Open Virtualization Format (OVF and OVA) and XenServer XVA format.

You can also use the xe CLI to import VMs from and export them to XenServer XVA format.

Requirements

To be able to import and export VMs to and from different formats, you require certain permissions. The **Import** or **Export** wizard performs checks to ensure that you have a required role in the destination pool before allowing you to continue. For more information about the RBAC roles that are required, see [RBAC roles and permissions](#).

Supported formats

Format	Supported for import?	Supported for export?	Description
Open Virtualization Format (OVF and OVA)	Yes	Yes	OVF is an open standard for packaging and distributing a virtual appliance consisting of one or more VMs.
Disk image formats (VHD, VHDX/AVHDX, and VMDK)	Yes	No	Virtual Hard Disk (VHD), Virtual Hard Disk Extended (VHDX), Automatic Virtual Hard Disk (AVHDX), and Virtual Machine Disk (VMDK) format disk image files can be imported using the Import wizard.
XenServer XVA format	Yes	Yes	XVA is a format specific to Xen-based hypervisors for packaging an individual VM as a single file archive, including a descriptor and disk images. Its file name extension is .xva .

Which format to use?

Consider using OVF/OVA format to:

- Share XenServer vApps and VMs with other virtualization platforms that support OVF
- Save more than one VM

Consider using VHD, VHDX/AVHDX or VMDK format when:

- There is a virtual disk image available without any associated OVF metadata

Consider using XVA format to:

- Import and export VMs by running a script on the CLI

Open virtualization format (OVF and OVA)

OVF is an open standard, specified by the Distributed Management Task Force, for packaging and distributing a virtual appliance consisting of one or more VMs. For further details about OVF and OVA formats, see the following information:

- [Open Virtualization Format Specification](#)

An **OVF Package** is the set of files that comprises the virtual appliance. It always includes a descriptor file and any other files that represent the following attributes of the package:

Attributes Descriptor (.ovf): The descriptor always specifies the virtual hardware requirements of the package. It may also specify other information, including:

- Descriptions of virtual disks, the package itself, and guest operating systems
- A license agreement
- Instructions to start and stop VMs in the appliance
- Instructions to install the package

Signature (.cert): The signature is the digital signature used by a public key certificate in the X.509 format to authenticate the author of the package.

Manifest (.mf): The manifest allows you to verify the integrity of the package contents. It contains the SHA-1 digests of every file in the package.

Virtual disks: OVF does not specify a disk image format. An OVF package includes files comprising virtual disks in the format defined by the virtualization product that exported the virtual disks. XenServer produces OVF packages with disk images in Dynamic VHD format; VMware products and Virtual Box produce OVF packages with virtual disks in Stream-Optimized VMDK format.

OVF packages also support other non-metadata related capabilities, such as compression, archiving, EULA attachment, and annotations.

Note:

When importing an OVF package that has been compressed or contains compressed files, you may need to free up extra disk space on the XenServer host to import it properly.

An **Open Virtual Appliance (OVA) package** is a single archive file, in the Tape Archive (.tar) format, containing the files that comprise an OVF Package.

Select OVF or OVA format OVF packages contain a series of uncompressed files, which makes it easier when you want to access individual disk images in the file. An OVA package contains one large file, and while you can compress this file, it does not give you the flexibility of a series of files.

Using the OVA format is useful for specific applications for which it is beneficial to have just one file, such as creating packages for Web downloads. Consider using OVA only as an option to make the package easier to handle. Using this format lengthens both the export and import processes.

Disk image formats (VHD, VHDX/AVHDX, and VMDK)

Using XenCenter, you can import disk images in the Virtual Hard Disk (VHD), Virtual Hard Disk Extended (VHDX), Automatic Virtual Hard Disk (AVHDX), and Virtual Machine Disk (VMDK) formats. Exporting standalone disk images is not supported.

VHD is a virtual hard disk format that is used to store the disk data for VMs. VMDK is a virtual disk format that is specific to VMware. VHDX is an improved version of the VHD format. For more information about VHDX/AVHDX, see VHDX/AVHDX.

You might choose to import a disk image when a virtual disk image is available without any associated OVF metadata. This option might occur in the following situations:

- It is possible to import a disk image, but the associated OVF metadata is not readable
- A virtual disk is not defined in an OVF package
- You are moving from a platform that does not allow you to create an OVF package (for example, older platforms or images)
- You want to import an older VMware appliance that does not have any OVF information
- You want to import a standalone VM that does not have any OVF information

When available, we recommend importing appliance packages that contain OVF metadata rather than an individual disk image. The OVF data provides information that the **Import** wizard requires to recreate a VM from its disk image. This information includes the number of disk images associated with the

VM, processor, storage, network, and memory requirements. Without this information, it can be much more complex and error-prone to recreate the VM.

VHDX/AVHDX VHDX is an improved version of the VHD format that provides better performance, fault tolerance, and data protection features in addition to a larger capacity.

AVHDX is a disk image format that is part of the VHDX standard for virtual hard disks. It is used to store the snapshot or checkpoint information of VMs. Whenever you create a snapshot on a VM, an AVHDX file is automatically generated to preserve the current state of the virtual hard disk. This file can be used to track the snapshot information of VMs. Each AVHDX file is linked to the previous AVHDX file, creating a chain that contains all changes made since the base virtual hard disk file.

Using the XenCenter **Import** wizard, you can import a VHDX/AVHDX file into a pool or specific host as a VM. For information on how to use the XenCenter **Import** wizard to import a VHDX/AVHDX file, see [Import disk images](#).

If you are a Citrix customer, you can also use this feature to migrate your VMs from VMware to XenServer. As part of its software streaming technology, Citrix Provisioning stores a shared disk image (vDisk) as a VHDX/AVHDX file. By using the XenCenter **Import** wizard to import a VHDX/AVHDX file, you can easily migrate your vDisk from VMware to XenServer. For more information, see [Migrating from VMware](#).

XVA format

XVA is a virtual appliance format specific to XenServer, which packages a single VM as a single set of files, including a descriptor and disk images. The file name extension is **.xva**.

The descriptor (file name extension **ova.xml**) specifies the virtual hardware of a single VM.

The disk image format is a directory of files. The directory name corresponds to a reference name in the descriptor and contains two files for each 1 MB block of the disk image. The base name of each file is the block number in decimal. The first file contains one block of the disk image in raw binary format and does not have an extension. The second file is a checksum of the first file. If the VM was exported from Citrix Hypervisor 8.0 or earlier, this file has the extension **.checksum**. If the VM was exported from Citrix Hypervisor 8.1 or later, this file has the extension **.xxhash**.

Important:

If a VM is exported from the XenServer host and then imported into another XenServer host with a different CPU type, it may not run properly. For example, a Windows VM exported from a host with an Intel VT Enabled CPU might not run when imported into a host with an AMD-VTM CPU.

Operating system fixup

When importing a virtual appliance or disk image created and exported from a virtualization platform other than XenServer, you might have to configure the Linux VM before it boots properly on the XenServer host.

XenCenter includes an advanced hypervisor interoperability feature—Operating System Fixup—which aims to ensure a basic level of interoperability for VMs that you import into XenServer. Use Operating System Fixup when importing VMs from OVF/OVA packages and disk images created on other virtualization platforms.

The Operating System Fixup ISO is located in your XenCenter installation in the directory `C:\Program Files (x86)\XenServer\XenCenter\External Tools`.

The Operating System Fixup process addresses the operating system device and driver issues inherent when moving from one hypervisor to another. The process attempts to repair boot device-related problems with the imported VM that might prevent the operating system within from booting in the XenServer environment. This feature is not designed to perform conversions from one platform to another.

Note:

This feature requires an ISO storage repository with 40 MB of free space and 256 MB of virtual memory.

Operating System Fixup is supplied as an automatically booting ISO image that is attached to the DVD drive of the imported VM. It performs the necessary repair operations when the VM is first started, and then shuts down the VM. The next time the new VM is started, the boot device is reset, and the VM starts normally.

To use Operating System Fixup on imported disk images or OVF/OVA packages, enable the feature on the **OS Fixup Settings** page of the XenCenter **Import** wizard. Specify a location where the Fixup ISO is copied so that XenServer can use it.

For Windows VMs, if your pool has Conversion Manager 8.5.0 or later installed in it, the import process calls the Conversion Manager to prepare the imported Windows VMs for use. The operating system fixup option is not available in this case.

What does operating system fixup do to the VM?

The Operating System Fixup option is designed to make the minimal changes possible to enable a virtual system to boot. Depending on the guest operating system and the hypervisor of the original host, further actions might be required after using Operating System Fixup. These actions can include configuration changes and driver installation.

During the Fixup process, an ISO is copied to an ISO SR. The ISO is attached to a VM. The boot order is set to boot from the virtual DVD drive, and the VM boots into the ISO. The environment within the ISO then checks each disk of the VM to determine if it is a Linux or a Windows system.

If a Linux system is detected, the location of the GRUB configuration file is determined. Any pointers to SCSI disk boot devices are modified to point to IDE disks. For example, if GRUB contains an entry of `/dev/sda1` representing the first disk on the first SCSI controller, this entry is changed to `/dev/hda1` representing the first disk on the first IDE controller.

If certain virtualization tool sets are discovered in the VM, they are disabled to prevent performance problems and unnecessary event messages.

Note:

The Operating System Fixup option is not supported for the RHEL family of distributions. When booted, the imported RHEL VM will enter emergency mode. To work around this problem, reboot it into rescue mode, run `dracut --regenerate-all --force`, and then reboot the VM again.

Import VMs

When you import a VM, you effectively create a VM, using many of the same steps required to provision a new VM. These steps include nominating a host, and configuring storage and networking.

You can import OVF/OVA, disk image, XVA, and XVA Version 1 files using the XenCenter **Import** wizard. You can also import XVA files via the xe CLI.

Note:

It is not guaranteed that you can import a VM that was exported from a more up-to-date system onto a less up-to-date system. For example, if you export a VM from a host that is at the latest level on the Early Access update channel, you might not be able to import this VM into a host that is at the latest level on the Normal update channel straight away. Before you can import the VM into that host, the host on the Normal channel must reach the same level of updates as the Early Access channel had when the VM was exported.

This behavior is caused by additional metadata that might be added in later updates, which earlier versions do not understand.

Import VMs from OVF/OVA

The XenCenter **Import** wizard allows you to import VMs that have been saved as OVF/OVA files. The **Import** wizard takes you through the usual steps to create a VM in XenCenter: nominating a host, and

then configuring storage and networking for the new VM. When importing OVF and OVA files, extra steps may be required, such as:

- When importing Linux VMs that have been created using other virtualization platforms, run the Operating System Fixup feature to ensure a basic level of interoperability for the VM. For more information, see [Operating system fixup](#).

Prerequisites

- Ensure that the target host has enough RAM to support the virtual machines being imported. A lack of available RAM results in a failed import.

Imported OVF packages appear as vApps when imported using XenCenter. When the import is complete, the new VMs appear in the XenCenter **Resources** pane, and the new vApp appears in the **Manage vApps** dialog box.

To import VMs from OVF/OVA by using XenCenter:

1. In the **Resources** pane, select the host or pool to import the package into.
2. On the **File** menu, select **Import**. The **Import** wizard opens.
3. On the first page of the wizard, locate the OVA/OVF file you want to import, then click **Next** to continue.
 - If you select a compressed OVA file (*.ova.gz), on clicking **Next**, the file is decompressed to an OVA file and the old *.ova.gz file is deleted.
 - If you enter a URL location (HTTP, HTTPS, file, FTP) in the **Filename** box, on clicking **Next**, a **Download Package** dialog opens. In this dialog, specify a folder on your XenCenter host where the OVA/OVF is to be copied.
4. (Optional) **Review and accept the EULA**. If the OVF/OVA includes a EULA, the **EULA** page displays the EULA text. If you want to accept the EULA, choose **Accept** and continue.
5. **Specify a template**. In the **Template** page, choose the template to use for the VM or VMs that you are importing from the OVA/OVF file. Only the default templates provided with XenServer are supported. Do not use a custom template.

On this panel, you can also choose whether to use Windows Update to update the I/O drivers. Click **Next** to continue.
6. (Optional) **Install the XenServer VM Tools** (Windows VMs only) If the Conversion Manager virtual appliance version 8.5.0 or later is running in your pool, you can choose to install XenServer VM Tools for Windows as part of the import. This tools installation process happens after the imported VMs first start. For this option to be available, the VM disks must be on shared storage or on the same host as the Conversion Manager.

- a) Select **Install the XenServer® VM Tools**.
 - b) Choose which version of the tools to install. One or more of the following options are available:
 - The version of the tools uploaded into the Conversion Manager
 - The latest version downloaded from the internet
 - A tools installation file in a location accessible from XenCenter®
 - c) Click **Next** to continue.
7. (Optional) **Specify the VM location.** On the **Location** page, choose a server if you want to place the VMs you are importing on the server's local storage.
- Click **Next** to continue.
8. **Configure storage for the new VMs.**
- On the **Storage** page, select a storage repository (SR) where the imported virtual disk is placed. If you have selected to install the XenServer VM Tools as part of the import, ensure that VM disks are on shared storage or on the same host as the Conversion Manager.
- Click **Next** to continue.
9. **Configure networking for the new VMs.**
- On the **Networking** page, select a target network in the destination pool/standalone server for the new VMs' virtual network interface.
- Click **Next** to continue.
10. (Optional) **Specify security settings.** If the selected OVF/OVA package is configured with security features such as certificates or a manifest, specify the necessary information on the **Security** page. Select **Next** to continue. Different options appear on this page depending on which security features have been configured on the OVF package:
- If the package is signed, a **Verify digital signature** check box appears here; select this check box if you want to verify the signature. Click **View Certificate** to display the certificate used to sign the package. If the certificate appears as untrusted, it is likely that either the Root Certificate or the Issuing Certificate Authority is not trusted on the local computer.
 - If the package includes a manifest, a **Verify manifest content** check box appears here. Select this check box to have the wizard verify the list of files in the package.
- When the package is digitally signed, the associated manifest is verified automatically and so the **Verify manifest content** check box does not appear on the **Security** page.
- Click **Next** to continue.

11. (Optional) **Enable Operating System Fixup.** (Linux VMs only) If all your VMs are Windows VMs, this panel is not shown.

If the Linux VMs that you are importing are built on a hypervisor other than XenServer, select **Use Operating System Fixup**. Choose an ISO SR where the Fixup ISO can be copied so that XenServer can use it. For more information, see [About VM Import and Export](#).

Click **Next** to continue.

12. On the **Finish** page, review all the import settings. Click **Finish** to begin the import process and close the wizard.

When the import is finished, the newly imported VMs appear in the **Resources** pane and the new vApp appears in the [Managing vApps](#) dialog.

Note:

If Conversion Manager virtual appliance version 8.5.0 or later is running in your pool, it automatically removes the VMware tools, if they are present, after the imported Windows VMs first start.

Import disk images

The XenCenter **Import** wizard allows you to import a disk image into a pool or specific host as a VM. The **Import** wizard takes you through the usual steps to create a VM in XenCenter: nominating a host, and then configuring storage and networking for the new VM.

Prerequisites

- Ensure that DHCP runs on the management network XenServer is using.
- The **Import** wizard requires local storage on the server on which you are running it.

To import VMs from a disk image by using XenCenter

1. In the **Resources** pane, select the host or pool to import the package into.
2. On the **File** menu, select **Import**. The **Import** wizard opens.
3. On the first page of the wizard, locate the disk image file you want to import, then click **Next** to continue.

If you enter a URL location (HTTP, HTTPS, file, FTP) in the **Filename** box, on clicking **Next**, a **Download Package** dialog opens. In this dialog, specify a folder on your XenCenter host where the disk image is to be copied.

4. **Specify a template.** Choose the template to use for the VM that you are importing the disk image for. Only the default templates provided with XenServer are supported. Do not use a custom template.

On this panel, you can also choose whether to copy the BIOS strings to the VM and whether to use Windows Update to update the I/O drivers.

Click **Next** to continue.

5. (Optional) **Install the XenServer VM Tools** (Windows VMs only) If the Conversion Manager virtual appliance version 8.5.0 or later is running in your pool, you can choose to install XenServer VM Tools for Windows as part of the import. This tools installation process happens after the imported VM first starts. For this option to be available, the disk image must be on shared storage or on the same host as the Conversion Manager.

- a) Select **Install the XenServer VM Tools**.

- b) Choose which version of the tools to install. One or more of the following options are available:

- The version of the tools uploaded into the Conversion Manager
- The latest version downloaded from the internet
- A tools installation file in a location accessible from XenCenter

- c) Click **Next** to continue.

6. **Specify the VM name.**

On the **Name** page, enter the name of the new VM to be created from the imported disk image. This value can be adjusted later, after the new VM has been created. VM names are not checked for uniqueness within XenCenter, so it makes it easier for you to manage different VMs if you give them meaningful, memorable names. For more information, see [VM Name and Description](#).

Click **Next** to continue.

7. **Allocate CPU and memory resources.**

On the **CPU & Memory** page, allocate CPU and initial memory resources. These values can be adjusted later, after the new VM has been created. For more information, see [VM CPU and Memory Allocation](#).

Click **Next** to continue.

8. (Optional) **Specify the VM location.** On the **Location** page, choose a server if you want to place the VM you are importing on the server's local storage.

Click **Next** to continue.

9. **Configure storage for the new VM.**

On the **Storage** page, select a storage repository (SR) where the imported virtual disk is placed. If you have selected to install the XenServer VM Tools as part of the import, ensure that disk image is on shared storage or on the same host as the Conversion Manager.

Click **Next** to continue.

10. **Configure networking for the new VM.**

On the **Networking** page, select a target network in the destination pool/standalone server for the new VM's virtual network interface.

Click **Next** to continue.

11. **Specify the boot option.**

On the **Boot options** page, specify the boot mode for the new VM. Only the options that are supported with the guest operating system can be selected.

- **BIOS Boot:** Select this option to boot the VM in legacy BIOS mode. Legacy BIOS boot is not supported for Windows 11 VMs.
- **UEFI Boot:** Select this option to boot the VM in UEFI mode. XenServer supports this boot mode only for certain guest operating systems. For more information about the guest operating systems that are supported for UEFI boot, see [Guest UEFI boot and Secure Boot](#).
- **UEFI Secure Boot:** Select this option to boot the VM in UEFI Secure Boot mode. Secure Boot prevents unsigned, incorrectly signed or modified binaries from being run during boot. XenServer supports this boot mode only for certain guest operating systems. For more information about the guest operating systems that are supported for UEFI Secure Boot, see [Guest UEFI boot and Secure Boot](#).

(VHDX/AVHDX only) When importing a VM from a VHDX/AVHDX file, the **Import** wizard attempts to automatically detect the boot mode from the VHDX/AVHDX file you are importing. Review the settings and ensure that the same boot mode that was originally configured for the VM is selected.

Click **Next** to continue.

12. (Optional) **Enable Operating System Fixup.** If you have selected a Windows template, this panel is not shown.

If the disk image that you are importing is built on a hypervisor other than XenServer, select **Use Operating System Fixup**. Choose an ISO SR where the Fixup ISO can be copied so that XenServer can use it. For more information, see [About VM Import and Export](#).

Important:

If you are importing a VHDX/AVHDX file from Citrix Provisioning (PVS) on VMware vSphere, select **Don't use Operating System Fixup**. Enabling Operating System Fixup for

VHDX/AVHDX from PVS might result in a boot failure.

Click **Next** to continue.

13. On the **Finish** page, review all the import settings. Click **Finish** to begin the import process and close the wizard.

The import progress is displayed on the status bar at the bottom of the XenCenter window and also on the **Events** view under **Notifications**.

The import process might take some time. The time it takes depends on the size of the imported virtual disks, the available network bandwidth, and the disk interface speed of the XenCenter host. When the import is finished, the newly imported VM appears in the **Resources** pane.

Note:

If Conversion Manager virtual appliance version 8.5.0 or later is running in your pool, it automatically removes the VMware tools, if they are present, after the imported VM first starts.

Next steps Ensure the XenServer VM Tools for Windows are fully installed:

- If you chose to install XenServer VM Tools for Windows as part of the import process, restart the VM after the import to complete the tools installation.
- If you chose not to install XenServer VM Tools for Windows as part of the import process, ensure that you install them manually later for the VM to be fully supported. For more information, see [Install XenServer VM Tools](#).

Import VMs from XVA

You can import VMs, templates, and snapshots that have previously been exported and stored locally in XVA format (.xva). To do so, you follow the usual steps to create a VM: nominating a host, and then configuring storage and networking for the new VM.

Warning:

It may not always be possible to run an imported VM that was exported from another host with a different CPU type. For example, a Windows VM exported from a host with an Intel VT Enabled CPU might not run when imported to a host with an AMD-VTM CPU.

To import VMs from XVA by using XenCenter:

1. Open the **Import** wizard by doing one of the following:
 - In the **Resources** pane, right-click, and then select **Import** on the shortcut menu.
 - On the **File** menu, select **Import**.

2. On the first page of the wizard, locate the file you want to import (.xva or ova.xml), and then click **Next** to continue.

If you enter a URL location (http, https, file, or ftp) in the **Filename** box. Click **Next**, a Download Package dialog box opens and you must specify a folder on your XenCenter host where the file is copied.

3. Select a pool or host for the imported VM to start on, and then choose **Next** to continue.
4. Select the storage repositories on which to place the imported virtual disk, and then click **Next** to continue.
5. Configure networking for the imported VMs: map the virtual network interface in the VM you are importing to target a network in the destination pool. The Network and MAC address shown in the list of incoming VMs are stored as part of the definition of the original (exported) VM in the export file. To map an incoming virtual network interface to a target network, select a network from the list in the Target Network column. Click **Next** to continue.
6. Review the import settings, and then click **Finish** to begin the import process and close the wizard.

Note:

Importing a VM may take some time, depending on the size of the VM and the speed and bandwidth of the network connection.

The import progress is displayed in the status bar at the bottom of the XenCenter window and on the **Notifications > Events** view. When the newly imported VM is available, it appears in the **Resources** pane.

To import a VM from XVA by using the xe CLI:

To import the VM to the default SR on the target XenServer host, enter the following:

```
1 xe vm-import -h hostname -u root -pw password \  
2   filename=pathname_of_export_file
```

To import the VM to a different SR on the target XenServer host, add the optional `sr-uuid` parameter:

```
1 xe vm-import -h hostname -u root -pw password \  
2   filename=pathname_of_export_file sr-uuid=uuid_of_target_sr
```

If you want to preserve the MAC address of the original VM, add the optional `preserve` parameter and set to `true`:

```
1 xe vm-import -h hostname -u root -pw password \  
2   filename=pathname_of_export_file preserve=true
```


Note:

Importing a VM may take some time, depending on the size of the VM and the speed and bandwidth of the network connection.

After the VM has been imported, the command prompt returns the UUID of the newly imported VM.

Export VMs

You can export OVF/OVA and XVA files using the XenCenter Export wizard. You can also export XVA files by using the xe CLI.

Note

VM group settings are not retained when a VM is exported.

Export VMs as OVF/OVA

Using the XenCenter Export wizard, you can export one or more VMs as an OVF/OVA package. When you export VMs as an OVF/OVA package, the configuration data is exported along with the virtual hard disks of each VM.

To export VMs as OVF/OVA by using XenCenter:

1. Shut down or suspend the VMs that you want to export.
2. Open the Export wizard: in the **Resources** pane, right-click the pool or host containing the VMs you want to export, and then select **Export**.
3. On the first page of the wizard:
 - Enter the name of the export file
 - Specify the folder where you want the files to be saved
 - Select **OVF/OVA Package (*.ovf, *.ova)** from the **Format** list
 - Click **Next** to continue
4. From the list of available VMs, select the VMs that you want to include in the OVF/OVA package, and then click **Next** to continue.
5. If necessary, you can add to a previously prepared End User Licensing Agreement (EULA) document (.rtf, .txt) to the package.

To add a EULA, click **Add** and browse to the file you want to add. Once you have added the file, you can view the document by selecting it from the **EULA files** list and then clicking **View**.

EULAs can provide the legal terms and conditions for using the appliance and the applications delivered in the appliance.

The ability to include one or more EULAs lets you legally protect the software on the appliance. For example, if your appliance includes a proprietary operating system on its VMs, you might want to include the EULA text from that operating system. The text is displayed and the person who imports the appliance must accept it.

Note:

Attempting to add EULA files that are not in supported formats, including XML or binary files, can cause the import EULA functionality to fail.

Select **Next** to continue.

6. On the **Advanced options** page, specify a manifest, signature and output file options, or just click **Next** to continue.

- a) To create a manifest for the package, select the **Create a manifest** check box.

The manifest provides an inventory or list of the other files in a package. The manifest is used to ensure that the files originally included when the package was created are the same files present when the package arrives. When the files are imported, a checksum is used to verify that the files have not changed since the package was created.

- b) To add a digital signature to the package

- i. Select **Sign the OVF package**.

The digital signature ([.cert](#)) contains the signature of the manifest file and the certificate used to create that signature. When a signed package is imported, the user can verify the identity of the package creator by using the public key of the certificate to validate the digital signature.

- ii. Browse to locate a certificate.

Use an X.509 certificate that you have already created from a Trusted Authority and exported as a [.pfx](#) file. For certificates with SHA-256 digest export using the “Microsoft Enhanced RSA and AES Cryptographic Provider” as CSP.

- iii. In **Private key password** enter the export (PFX) password, or, if an export password was not provided, the private key associated with the certificate.

- c) To output the selected VMs as a single (tar) file in OVA format, select the **Create OVA package (single OVA export file)** check box. For more on the different file formats, see [Open virtualization format](#).
- d) To compress virtual hard disk images (.VHD files) included in the package, select the **Compress OVF files** check box.

When you create an OVF package, the virtual hard disk images are, by default, allocated the same amount of space as the exported VM. For example, a VM that is allocated 26 GB of space has a hard disk image that consumes 26 GB of space. The hard disk image uses this space regardless of whether or not the VM actually requires it.

Note:

Compressing the VHD files makes the export process take longer to complete. Importing a package containing compressed VHD files also takes longer, as the **Import** wizard must extract all of the VHD images as it imports them.

If both **Create OVA package (single OVA export file)** and **Compress OVF files** are checked, the result is a compressed OVA file with the extension `.ova.gz`.

7. Review the export settings.

To have the wizard verify the exported package, select the **Verify export on completion** check box. Click **Finish** to begin the export process and close the wizard.

Note:

Exporting a VM may take some time, depending on the size of the VM and the speed and bandwidth of the network connection.

The export progress is displayed in the status bar at the bottom of the XenCenter window and on the **Notifications > Events** view. To cancel an export in progress, click the **Notifications > Events** view, find the export in the list of events, and click the **Cancel** button.

Export VMs as XVA You can export an existing VM as an XVA file using the XenCenter **Export** wizard or the xe CLI. We recommend exporting a VM to a machine other than the XenServer host, on which you can maintain a library of export files. For example, you can export the VM to the machine running XenCenter.

Warning:

It may not always be possible to run an imported VM that was exported from another host with a different CPU type. For example, a Windows VM exported from a host with an Intel VT Enabled CPU might not run when imported to a host with an AMD-VTM CPU.

To export VMs as XVA files by using XenCenter:

1. Shut down or suspend the VM that you want to export.
2. Open the Export wizard: from the **Resources** pane, right-click the VM which you want to export, and then select **Export**.

3. On the first page of the wizard:

- Enter the name of the export file
- Specify the folder where you want the files to be saved
- Select **XVA File (*.xva)** from the **Format** list
- Click **Next** to continue

4. From the list of available VMs, select the VM that you want to export, and then click **Next** to continue.

5. Review the export settings.

To have the wizard verify the exported package, select the **Verify export on completion** check box. Click Finish to begin the export process and close the wizard.

Note:

Exporting a VM may take some time, depending on the size of the VM and the speed and bandwidth of the network connection.

The export progress is displayed in the status bar at the bottom of the XenCenter window and on the **Notifications > Events** view. To cancel an export in progress, click the **Notifications > Events** view, find the export in the list of events, and click the **Cancel** button.

To export VMs as XVA files by using the xe CLI:

1. Shut down the VM that you want to export.
2. Export the VM by running the following:

```
1 xe vm-export -h hostname -u root -pw password vm=vm_name \  
2   filename=pathname_of_file
```

Note:

Be sure to include the **.xva** extension when specifying the export file name. If the exported VM doesn't have this extension, XenCenter might fail to recognize the file as a valid XVA file when you attempt to import it.

VM placement

December 17, 2025

XenServer® provides several ways that you can define which host in a pool your VMs are started on or migrated to:

- **Anti-affinity placement groups:** When VMs in a VM group with anti-affinity placement rules are started or evacuated from a host, XenServer places the VM on a host that enables the group to most closely follow the group rules. For more information, see [Anti-affinity placement groups](#).
- **Home server:** This VM setting defines the host where XenServer attempts to start the VM. If a home server is set, this setting has higher priority than any VM group placement rules on the same VM. For more information, see [Home server](#).
- **Workload Balancing:** A virtual appliance that monitors your pool and provides rebalancing and placement recommendations to maximize VM density or resource performance. By default, Workload Balancing placements take priority over the home server and anti-affinity placement group settings. Workload Balancing placement recommendations can breach the anti-affinity rules. For more information, see [Workload Balancing](#).

Anti-affinity placement groups

You can assign your VMs to VM groups that are started on hosts in your pool according to the placement rules of the group. The available placement types are “anti-affinity” and “normal”.

VM groups with normal placement do not have any special behavior.

VM groups with anti-affinity placement have the following rules:

1. No single host is the only place all VMs in a particular group are running.
2. The count of VMs within the group for each host is as even as possible.

When a VM in a VM group with anti-affinity placement rules is started or evacuated from a host, XenServer® chooses to place the VM on a host that enables the VM group to follow these anti-affinity placement rules as closely as possible. If the VM start or host evacuate cannot be done without breaching the placement rules, XenServer prioritizes completing the action over following the placement rules.

If the anti-affinity rules are breached, you receive an alert. Alerts are generated if all the running VMs in an anti-affinity group reside on a single host. If this changes, the alert is cleared. This alert is not generated if only one of the VMs in the group is running or no VMs in the group are running. For more information, see [XenCenter Alerts](#).

In the following scenarios the VM group placement rules do not take effect:

- In pools with Workload Balancing (WLB) enabled, the VM group is not used for starting, restarting, resuming, or migrating the VM. Instead, Workload Balancing nominates the best server for the VM by analyzing XenServer resource pool metrics and by recommending optimizations. These recommendations can cause the VM anti-affinity rules to be breached.
- If the VM has a home server defined, this setting takes priority over the VM group setting.

Considerations when using VM groups:

- Only 5 anti-affinity groups per pool are supported.
- When you clone a VM that is in a VM group, the new VM is not added to the same VM group. You must explicitly set the VM group for the new VM.
- If you revert a VM to a previous snapshot, the VM retains its current VM group setting. This setting does not change to the value it had when you took the snapshot.
- Exporting a VM does not retain the VM group setting. If you import this VM again, it is not added to a VM group. You must explicitly set the VM group for the imported VM.

Adding a VM to a group

You can add a VM to only one VM group.

In XenCenter®, use one the following methods to add a VM to a group:

- In the **New VM** wizard on the **VM group** tab, you can select a VM group to add the VM to or you can create a new one. For more information, see [Create a new VM](#).
- In the **Resources** pane of XenCenter, you can right-click on a VM or a selected set of VMs and choose **Assign to VM group**. The submenu shows the available VM groups and also gives you the option to create a new VM group. The current VM group for the VM is marked with a check.
- In the **Resources** pane of XenCenter, you can right-click on the pool and choose **Manage VM groups**. For more information, see [Manage VM groups](#).
- In the XenCenter main menu, you can select **VM > Assign to VM group**. The submenu shows the available VM groups and also gives you the option to create a new VM group. The current VM group for the VM is marked with a check.
- In the XenCenter main menu, you can select **Pool > Manage VM groups**. For more information, see [Manage VM groups](#).

To use the xe CLI to add a VM to a group, run the following command:

```
1 xe vm-param-set uuid=<vm_uuid> groups=<vm_group_uuid>
```

To get the UUIDs of the VM and the VM anti-affinity group, you can use the commands `xe vm-list` and `xe vm-group-list`.

Changing the group for a VM

A VM can be part of only one VM group. If you select to add a VM to a group, it is removed from the group it was previously a member of.

To change the group a VM is in:

In XenCenter, you can update the group for a VM in one of the following places:

- In the **Resources** pane of XenCenter, you can right-click on a VM and choose **Assign to VM group**. The submenu shows the available VM groups and also gives you the option to create a new VM group. The current VM group for the VM is marked with a check.
- In the **Resources** pane of XenCenter, you can right-click on the pool and choose **Manage VM groups**. For more information, see [Manage VM groups](#).
- In the XenCenter main menu, you can select **VM > Assign to VM group**. The submenu shows the available VM groups and also gives you the option to create a new VM group. The current VM group for the VM is marked with a check.
- In the XenCenter main menu, you can select **Pool > Manage VM groups**. For more information, see [Manage VM groups](#).

To use the xe CLI to change the group for a VM, run the following command:

```
1 xe vm-param-set uuid=<vm_uuid> groups=<vm_group_uuid>
```

To get the UUIDs of the VM and the VM anti-affinity group, you can use the commands `xe vm-list` and `xe vm-group-list`.

You can also use this command to remove the VM from an affinity group. To do this, leave the value of the `groups` parameter blank.

```
1 xe vm-param-set uuid=<vm_uuid> groups=
```

Working with VM groups

You can manage your VM groups in XenCenter by using the **Manage VM group** dialog box. For more information, see [Manage VM groups](#).

XenServer provides the following xe CLI commands for working with VM groups:

To create a VM group:

```
1 xe vm-group-create placement=anti-affinity name-label=<name_label>
```

Replace `<name-label>` with the name for your group. You can also provide an optional `name-description` parameter that sets a description for the group. To create a group whose members follow the anti-affinity placement rules, specify `placement=anti-affinity`. The option `normal` is also available for the parameter `placement`. Normal placement groups have no special behavior or rules.

To see the available VM groups:

```
1 xe vm-group-list
```

To delete a VM group:

```
1 xe vm-group-destroy uuid=<vm_group_uuid>
```

You can list the available VM groups to find the group UUID.

If you delete a VM group that has VMs in it, the VMs are no longer associated with a group, but are not otherwise affected.

Home server

A home server is the server which provides the resources for a VM in a pool. When you nominate a home server for a VM, XenServer always attempts to start up the VM on that server if it can. If XenServer cannot start the VM on that server, an alternate server within the same pool is selected automatically.

In the following scenarios, the home server nomination does not take effect:

- In pools with Workload Balancing (WLB) enabled, the nominated home server is not used for starting, restarting, resuming, or migrating the VM. Instead, Workload Balancing nominates the best server for the VM by analyzing XenServer resource pool metrics and by recommending optimizations.
- If a VM has one or more virtual GPUs assigned to it, the home server nomination does not take effect. Instead, the server nomination is based on the virtual GPU placement policy set by the user. For more information, see [GPU Placement Policy](#).

Nominating the home server

To set a home server when first creating the VM:

In XenCenter:

- To nominate a home server for the VM in the **New VM** wizard, select **Place the VM on this server** and choose a server from the list.
- If you do not want to nominate a home server, click **Don't assign this VM a home server**. The VM uses the resources on the most suitable available server.

For more information, see [Create a new VM](#).

By using the xe CLI:

- To nominate a home server, during VM create, run the command `xe vm-param-set uuid=<vm_uuid> affinity=<host_uuid>`.

Note:

If you are creating a BIOS-customized VM, the OEM server from which you copy the BIOS strings is automatically selected as the home server for the new VM.

Changing the home server

To change the home server configuration for a VM:

In XenCenter:

- Go to the VM's **Properties** dialog box, then to the **Home Server** tab. For more information, see [Change VM properties](#).

By using the xe CLI:

- To change the home server, run the command `xe vm-param-set uuid=<vm_uuid> affinity=<host_uuid>`.
- To remove the home server from a VM, run the command `xe vm-param-set uuid=<vm_uuid> affinity=`. (Leave the value of `affinity` blank.)

Delete VMs

December 17, 2025

You can delete VMs by using the xe CLI or XenCenter.

Deleting a virtual machine (VM) removes its configuration and its filesystem from the host. When you delete a VM, you can choose to delete or preserve any virtual disks attached to the VM, in addition to any snapshots of the VM.

Delete a VM by using the xe CLI

To delete a VM:

1. Find the VM UUID:

```
1 xe vm-list
```

2. Shutdown the VM:

```
1 xe vm-shutdown uuid=<uuid>
```

3. (Optional) You can choose to delete the attached virtual disks:

a) Find the virtual disk UUIDs:

```
1 xe vm-disk-list vm=<uuid>
```

b) Delete the virtual disk:

```
1 xe vdi-destroy uuid=<uuid>
```

Important:

Any data stored in the VM's virtual disk drives is lost.

4. (Optional) You can choose to delete the snapshots associated with the VM:

a) Find the UUIDs of the snapshots:

```
1 xe snapshot-list snapshot-of=<uuid>
```

b) For each snapshot to delete, find the UUIDs of the virtual disks for that snapshot:

```
1 xe snapshot-disk-list snapshot-uuid=<uuid>
```

c) Delete each snapshot disk:

```
1 xe vdi-destroy uuid=<uuid>
```

d) Delete the snapshot:

```
1 xe snapshot-destroy uuid=<uuid>
```

5. Delete the VM:

```
1 xe vm-destroy uuid=<uuid>
```

Delete a VM by using XenCenter

To delete a VM:

1. Shut down the VM.
2. Select the stopped VM in the **Resources** panel, right-click, and select **Delete** on the shortcut menu. Alternatively, on the **VM** menu, select **Delete**.
3. To delete an attached virtual disk, select its check box.

Important:

Any data stored in the VM's virtual disk drives is lost.

4. To delete a snapshot of the VM, select its check box.
5. Click **Delete**.

When the delete operation is completed, the VM is removed from the **Resources** pane.

Note:

VM snapshots whose parent VM has been deleted (*orphan snapshots*) can still be accessed from the **Resources** pane. These snapshots can be exported, deleted, or used to create VMs and templates. To view snapshots in the **Resources** pane, select **Objects** in the Navigation pane and then expand the **Snapshots** group in the Resources pane.

vApps

December 17, 2025

A vApp is a logical group of one or more related Virtual Machines (VMs) which can be started up as a single entity. When a vApp is started, the VMs contained within the vApp start in a user-predefined order. This feature enables VMs which depend upon one another to be automatically sequenced. An administrator no longer has to manually sequence the startup of dependent VMs when a whole service requires restarting (for instance for a software update). The VMs within the vApp do not have to reside on one host and can be distributed within a pool using the normal rules.

The vApp feature is useful in the Disaster Recovery situation. You can group all VMs that are on the same Storage Repository or all VMs that relate to the same Service Level Agreement (SLA).

Note:

vApps can be created and changed using both XenCenter and the xe CLI. For information on working with vApps using the CLI, see [Command Line Interface](#).

Manage vApps in XenCenter®

The **Manage vApps** dialog box enables you to create, delete, change, start, and shut down vApps, and import and export vApps within the selected pool. If you select a vApp in the list, the VMs it contains are listed in the details pane on the right.

You can use **Manage vApps** to do the following actions:

- To change the name or description of a vApp
- To add or remove VMs from the vApp
- To change the startup sequence of the VMs in the vApp

To change vApps:

1. Select the pool and, on the **Pool** menu, select **Manage vApps**.
Alternatively, right-click in the **Resources** pane and select **Manage vApps** on the shortcut menu.
2. Select the vApp and choose **Properties** to open its Properties dialog box.
3. Select the **General** tab to change the vApp name or description.
4. Select the **Virtual Machines** tab to add or remove VMs from the vApp.
5. Select the **VM Startup Sequence** tab to change the start order and delay interval values for individual VMs in the vApp.
6. Click **OK** to save your changes and close **Properties**.

Create vApps**To group VMs together in a vApp follow the procedure:**

1. Choose the pool and, on the **Pool** menu, select **Manage vApps**.
2. Type a name for the vApp, and optionally a description. Click **Next**.
You can choose any name you like, but a name that describes the vApp is best. Although it is advisable to avoid creating multiple vApps that have the same name, it is not a requirement. XenCenter does not force vApp names to be unique. It is not necessary to use quotation marks for names that include spaces.
3. Choose which VMs to include in the new vApp. Click **Next**.
You can use the search field to list only VMs that have names that include the specified text string.
4. Specify the startup sequence for the VMs in the vApp. Click **Next**.

Value	Description
Start Order	Specifies the order in which individual VMs are started up within the vApp, allowing certain VMs to be restarted before others. VMs that have a start order value of 0 (zero) are started first. VMs that have a start order value of 1 are started next. Then VMs that have a start order value of 2 are started, and so on.
Attempt to start next VM after	Specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence. That next group is the set of VMs that have a lower start order.

1. On the final page of **Manage vApps**, you can review the vApp configuration. Click **Previous** to go back and change any settings or **Finish** to create the vApp and close **Manage vApps**.

Note:

A vApp can span across multiple hosts in a single pool, but cannot span across several pools.

Delete vApps

To delete a vApp, follow the procedure:

1. Choose the pool and, on the **Pool** menu, select **Manage vApps**.
2. Select the vApp you want to delete from the list. Click **Delete**.

Note:

The VMs in the vApp are **not** deleted.

Start and shut down vApps by using XenCenter

To start or shut down a vApp, use **Manage vApps**, accessed from the **Pool** menu. When you start a vApp, all the VMs within it are started up automatically in sequence. The start order and delay interval values specified for each individual VM control the startup sequence. These values can be set when you first create the vApp. Change these values at any time from the vApp Properties dialog box or individual VM Properties dialog box.

To start a vApp:

1. Open **Manage vApps**: Choose the pool where the VMs in the vApp are located and, on the **Pool** menu, select **Manage vApps**. Alternatively, right-click in the **Resources** pane and select **Manage vApps** on the shortcut menu.
2. Choose the vApp and click **Start** to start all the VMs it contains.

To shut down a vApp:

1. Open **Manage vApps**: Choose the pool where the VMs in the vApp are located and, on the **Pool** menu, select **Manage vApps**. Alternatively, right-click in the **Resources** pane and select **Manage vApps** on the shortcut menu.
2. Choose the vApp and click **Shut Down** to shut down all the VMs in the vApp.

A soft shutdown is attempted on all VMs. If a soft shutdown is not possible, then a forced shutdown is performed.

Note:

A soft shutdown performs a graceful shutdown of the VM, and all running processes are halted individually.

A forced shutdown performs a hard shutdown and is the equivalent of unplugging a physical server. It might not always shut down all running processes. If you shut down a VM in this way, you risk losing data. Only use a forced shutdown when a soft shutdown is not possible.

Import and export vApps

vApps can be imported and exported as OVF/OVA packages. For more information, see [Import and Export VMs](#).

To export a vApp:

1. Open **Manage vApps**: on the **Pool** menu, select **Manage vApps**.
2. Choose the vApp you want to export in the list. Click **Export**.
3. Follow the procedure described in [Export VMs as OVF/OVA](#).

Exporting a vApp can take some time.

To import a vApp:

1. Open **Manage vApps**: on the **Pool** menu, select **Manage vApps**.
2. Click **Import** to open the **Import** dialog box.
3. Follow the procedure described in [Import VMs as OVF/OVA](#).

After the import is complete, the new vApp appears in the list of vApps in **Manage vApps**.

Manage UEFI Secure Boot certificate remediation

June 26, 2026

Microsoft's original 2011 UEFI Secure Boot certificates begin to expire in June 2026, in favor of replacement certificates issued in 2023. XenServer® provides a managed remediation workflow that helps administrators identify virtual machines (VMs) whose UEFI NVRAM only contains the legacy 2011 Microsoft certificate chain and safely update those VMs to include the 2023 certificates during a controlled reboot.

Virtual machines newly provisioned on XenServer 9 include the newer 2023 Microsoft certificates. This remediation workflow applies only to VMs originally provisioned on older XenServer releases, where they may only have the legacy 2011 certificates.

The following VM-class objects are in scope for this remediation workflow:

- UEFI Windows and Linux VMs that contain only the Microsoft 2011 UEFI certificate chain, including VMs where Secure Boot is currently disabled.
- Snapshots and templates that carry the legacy certificate state (exposed for reporting, not directly updatable).

The following are out of scope:

- VMs that already contain the 2023 Microsoft certificates.
- BIOS boot VMs.
- VMs using custom UEFI certificates.

Note that where virtual machines are being provisioned by Citrix Virtual Apps and Desktops™ (CVAD), additional steps may be required to ensure future virtual machines contain the 2023 certificates. Please refer to CVAD guidance in [CTX696696](#) for specific actions required.

Note:

The functionality described below was released as an update to the XenServer 9 Preview Early Access channel on [Jun 2, 2026](#), and the Normal channel on [Jun 8, 2026](#).

Certificate states

XenServer exposes a read-only certificate state for each VM, snapshot, and template object. The possible values are described in the following table:

State	Meaning
ok	The object does not require remediation. This includes BIOS boot VMs and VMs already carrying the 2023 Microsoft certificate chain.
update_available	The VM is eligible for remediation and can be marked for update on the next boot.
update_on_boot	The VM has been marked and XenServer will update the certificates automatically at the next boot of that VM.

Note:

Snapshots and templates display their certificate state for visibility only. Remediation applies only after a snapshot or template is converted to a VM and that VM is booted.

Prerequisites**Warning:**

Updating the UEFI certificates changes TPM PCR measurements (notably PCR7), which can trigger BitLocker recovery mode on Windows guests that seal their Volume Master Key against PCR7. Plan and stage rollouts carefully to avoid unplanned user disruption.

Before you mark any VM for remediation, do the following:

- Ensure the BitLocker recovery keys are accessible as a precaution.
- Suspend BitLocker protection on affected Windows VMs prior to updating the certificates.
- Schedule the first post-remediation reboot in a monitored maintenance window.

Manage certificate remediation by using XenCenter®

See [Secure Boot Certificates](#).

Manage certificate remediation by using the CLI**Review certificate state**

Run the following command to list the certificate state for all VM-class objects, including snapshots and templates:


```
1 xe vm-list params=name=label,uuid,is-a-template,is-snapshot,secureboot-  
certificates-state
```

Mark a VM for remediation

Run the following command to mark a VM so that XenServer updates the certificates at the next boot of that VM:

```
1 xe vm-update-secureboot-certificates-on-boot uuid=<vm-uuid> mark=true
```

Cancel pending updates

Run the following command to cancel a pending certificate update before the VM reboots:

```
1 xe vm-update-secureboot-certificates-on-boot uuid=<vm-uuid> mark=false
```

Note:

You can cancel a pending update only before the VM reboots and the scheduled remediation runs. After the VM has rebooted and the update has applied, the state returns to **ok** automatically.

XenAPI reference

The certificate remediation workflow introduces the following XenAPI field and method, which are first available in SDK version 26.14.0.

Field: `VM.secureboot_certificates_state`

Read-only. Applies to VM, snapshot, and template objects. Possible values: **ok**, **update_available**, **update_on_boot**.

Method: `VM.update_secureboot_certificates_on_boot(session, vm, mark)`

Pass **mark=true** to schedule the certificate update on the next boot. Pass **mark=false** to cancel a pending update.

PowerShell reference

With a VM object `$vm`, the new field can be accessed as follows:

```
1 $vm.secureboot_certificates_state
```

The new method is invoked as follows:

```
1 Invoke-XenVM -XenAction UpdateSecurebootCertificatesOnBoot -VM $vm -  
   Mark $true
```

Planning

1. Schedule updates in reasonably sized batches to limit the scope of any impacts from issues such as BitLocker recovery.
2. Test updates on representative VMs first and validate guest boot behavior before broad rollout.
3. Do not bulk-reboot a large number of VMs simultaneously; stagger reboots to allow time to respond to unexpected BitLocker or boot issues.

Troubleshooting

BitLocker recovery prompt after reboot

If BitLocker is configured and not suspended on a Windows VM, it is likely to prompt for a BitLocker recovery key after the first reboot following certificate remediation.

Enter the BitLocker recovery key to unlock the drive. After the VM boots successfully, BitLocker resumes protecting the volume automatically. To avoid this behaviour, suspend BitLocker protection on the VM before the scheduled remediation reboot, and resume it after the reboot has completed.

VM remains in `update_available` state after reboot

If a VM still shows `update_available` after a reboot, the update may not have been applied. Verify that the VM was marked before it rebooted by checking the state with the xe CLI. If the state is `update_available` rather than `ok`, mark the VM again and reboot.

Snapshot or template shows `update_available`

Snapshots and templates cannot be directly remediated. To update the certificate state of a snapshot or template, convert it to a VM and then reboot that VM after marking it for update.

VM memory

December 17, 2025

When you create a VM, a fixed amount of memory is allocated to the VM. You can use Dynamic Memory Control (DMC) to improve the utilization of physical memory in your XenServer® environment. DMC is a memory management feature that enables dynamic reallocation of memory between VMs.

XenCenter provides a graphical display of memory usage in its **Memory** tab. For more information, see the [XenCenter documentation](#).

Dynamic Memory Control (DMC) provides the following benefits:

- You can add or delete memory without restarting the VMs, providing a seamless experience to the user.
- When hosts are full, DMC allows you to start more VMs on these hosts, reducing the amount of memory allocated to the running VMs proportionally.

What is Dynamic Memory Control (DMC)?

XenServer DMC works by automatically adjusting the memory of running VMs, keeping the amount of memory allocated to each VM between specified minimum and maximum memory values, guaranteeing performance, and permitting greater density of VMs per host.

Without DMC, when a host is full, starting additional VMs fail with “out of memory” errors. To reduce the existing VM memory allocation and make room for more VMs, edit each VM’s memory allocation and then restart the VM. When using DMC, XenServer attempts to reclaim memory by automatically reducing the current memory allocation of running VMs within their defined memory ranges. XenServer attempts to reclaim memory even when the host is full.

Notes:

Dynamic Memory Control is not supported with VMs that have a virtual GPU.

The concept of dynamic range

For each VM, the administrator can set a dynamic memory range. The dynamic memory range is the range within which memory can be added/removed from the VM without requiring a restart. When a VM is running, the administrator can adjust the dynamic range. XenServer always guarantees to keep the amount of memory allocated to the VM within the dynamic range. Therefore adjusting it while the VM is running may cause XenServer to adjust the amount of memory allocated to the VM. The most extreme case is where the administrator sets the dynamic min/max to the same value, forcing XenServer to ensure that this amount of memory is allocated to the VM. If new VMs are required to start on “full” hosts, running VMs have their memory ‘squeezed’ to start new ones. The required extra memory is obtained by squeezing the existing running VMs proportionally within their pre-defined dynamic ranges

DMC allows you to configure dynamic minimum and maximum memory levels –creating a Dynamic Memory Range (DMR) that the VM operates in.

- Dynamic Minimum Memory: A lower memory limit that you assign to the VM.
- Dynamic Higher Limit: An upper memory limit that you assign to the VM.

For example, if the Dynamic Minimum Memory was set at 512 MB and the Dynamic Maximum Memory was set at 1,024 MB, it gives the VM a Dynamic Memory Range (DMR) of 512–1024 MB, within which it operates. XenServer *guarantees* always to assign each VM memory within its specified DMR when using DMC.

The concept of static range

Many operating systems that XenServer supports do not fully ‘understand’ the notion of dynamically adding or deleting memory. As a result, XenServer must declare the maximum amount of memory that a VM is asked to consume at the time that it restarts. Declaring the maximum amount of memory allows the guest operating system to size its page tables and other memory management structures accordingly. This introduces the concept of a static memory range within XenServer. The static memory range cannot be adjusted when the VM is running. For a particular boot, the dynamic range is constrained such as to be always contained within this static range. The static minimum (the lower bound of the static range) protects the administrator and is set to the lowest amount of memory that the OS can run with XenServer.

Note:

We recommend that you do not change the static minimum level as the static minimum level is set at the supported level per operating system. See the memory constraints table for more details.

Setting a static maximum level higher than a dynamic max allows you to allocate more memory to a VM in future without restarting the VM.

DMC behavior

Automatic VM squeezing

- If DMC is not enabled, when hosts are full, new VM starts fail with ‘out of memory’ errors.
- When DMC is enabled, even when hosts are full, XenServer attempts to reclaim memory by reducing the memory allocation of running VMs within their defined dynamic ranges. In this way, running VMs are squeezed proportionally at the same distance between the dynamic minimum and dynamic maximum for all VMs on the host

When DMC is enabled

- When the host's memory is plentiful - All running VMs receive their Dynamic Maximum Memory level
- When the host's memory is scarce - All running VMs receive their Dynamic Minimum Memory level.

When you are configuring DMC, remember that allocating only a small amount of memory to a VM can negatively impact it. For example, allocating too little memory:

- Using Dynamic Memory Control to reduce the amount of physical memory available to a VM can cause it to restart slowly. Likewise, if you allocate too little memory to a VM, it can start slowly.
- Setting the dynamic memory minimum for a VM too low can result in poor performance or stability problems when the VM is starting.

How does DMC work?

Using DMC, it is possible to operate a guest virtual machine in one of two modes:

1. **Target Mode:** The administrator specifies a memory target for the guest. XenServer adjusts the guest's memory allocation to meet the target. Specifying a target is useful in virtual server environments, and in situations where you know exactly how much memory you want a guest to use. XenServer adjusts the guest's memory allocation to meet the target you specify.
2. **Dynamic Range Mode:** The administrator specifies a dynamic memory range for the guest. XenServer selects a target from the range and adjusts the guest's memory allocation to meet the target. Specifying a dynamic range is useful in any situation where you want XenServer to repartition host memory dynamically in response to changing numbers of guests, or changing host memory pressure. XenServer selects a target from within the range and adjusts the guest's memory allocation to meet the target.

Note:

It is possible to change between target mode and dynamic range mode at any time for any running guest. Specify a new target, or a new dynamic range, and XenServer takes care of the rest.

Memory constraints

XenServer allows administrators to use all memory control operations with any guest operating system. However, XenServer enforces the following memory property ordering constraint for all guests:

`0 < memory-static-min <= memory-dynamic-min <= memory-dynamic-max <= memory-static-max`

XenServer allows administrators to change guest memory properties to any values that satisfy this constraint, subject to validation checks. However, in addition to the previous constraint, we support only certain guest memory configurations for each supported operating system. The range of supported configurations depends on the guest operating system in use. XenServer does not prevent administrators from configuring guests to exceed the supported limit. However, customers are advised to keep memory properties within the supported limits to avoid performance or stability problems. For detailed guidelines on the minimum and maximum memory limits for each supported operating system, see [Guest operating system support](#).

Warning:

When configuring guest memory, we advise NOT to exceed the maximum amount of physical memory addressable by your operating system. Setting a memory maximum that is greater than the operating system supported limit can lead to stability problems within your guest.

The dynamic minimum must be at least 75% of the static maximum for all supported operating systems. A lower limit can lead to stability problems. Administrators are encouraged to calibrate the sizes of their VMs carefully, and ensure that their working set of applications function reliably at dynamic-minimum.

xe CLI commands

Display the static memory properties of a VM

1. Find the UUID of the required VM:

```
1 xe vm-list
```

2. Note the uuid, and then run the command `param-name=memory-static`

```
1 xe vm-param-get uuid=uuid param-name=memory-static-{  
2   min,max }
```

For example, the following displays the static maximum memory properties for the VM with the UUID beginning ec77:

```
1 xe vm-param-get uuid= \  
2   ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 \  
3   param-name=memory-static-max;  
4   268435456
```

The example shows that the static maximum memory for this VM is 268,435,456 bytes (256 MB).

Display the dynamic memory properties of a VM

To display the dynamic memory properties, follow the procedure as above but use the command `param-name=memory-dynamic`:

1. Find the UUID of the required VM:

```
1 xe vm-list
```

2. Note the uuid, and then run the command `param-name=memory-dynamic`:

```
1 xe vm-param-get uuid=uuid param-name=memory-dynamic-{
2   min,max }
```

For example, the following displays the dynamic maximum memory properties for the VM with UUID beginning `ec77`

```
1 xe vm-param-get uuid= \
2   ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 \
3   param-name=memory-dynamic-max;
4   134217728
```

The example shows that the dynamic maximum memory for this VM is 134,217,728 bytes (128 MB).

Update memory properties

Warning:

Use the correct ordering when setting the static/dynamic minimum/maximum parameters. In addition, you must not invalidate the following constraint:

```
0 < memory-static-min <= memory-dynamic-min <= memory-dynamic-max
   <= memory-static-max
```

Update the static memory range of a virtual machine:

```
1 xe vm-memory-static-range-set uuid=uuid min=value max=value
```

Update the dynamic memory range of a virtual machine:

```
1 xe vm-memory-dynamic-range-set \
2   uuid=uuid min=value \
3   max=value
```

Specifying a target is useful in virtual server environments, and in any situation where you know exactly how much memory you want a guest to use. XenServer adjusts the guest's memory allocation to meet the target you specify. For example:

```
1 xe vm-memory-target-set target=value vm=vm-name
```

Update all memory limits (static and dynamic) of a virtual machine:

```
1 xe vm-memory-limits-set \  
2     uuid=uuid \  
3     static-min=value \  
4     dynamic-min=value \  
5     dynamic-max=value static-max=value
```

Notes:

- To allocate a specific amount memory to a VM that doesn't change, set the Dynamic Maximum and Dynamic Minimum to the same value.
- You cannot increase the dynamic memory of a VM beyond the static maximum.
- To alter the static maximum of a VM, you must shut down the VM.

Update individual memory properties**Warning:**

Do not change the static minimum level as it is set at the supported level per operating system. For more information, see [Memory constraints](#).

Update the dynamic memory properties of a VM.

1. Find the UUID of the required VM:

```
1 xe vm-list
```

2. Note the uuid, and then use the command `memory-dynamic-{ min,max } =value`

```
1 xe vm-param-set uuid=uuid memory-dynamic-{  
2     min,max }  
3     =value
```

The following example changes the dynamic maximum to 128 MB:

```
1 xe vm-param-set uuid=ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 memory-  
    dynamic-max=128MiB
```

Advanced notes for virtual machines

December 17, 2025

This section provides some advanced notes for Virtual Machines.

VM boot behavior

There are two options for the behavior of a Virtual Machine's VDI when the VM is booted:

Note:

The VM must be shut down before you can change its boot behavior setting.

Persist

Tip:

Use this boot behavior if you are hosting Citrix Virtual Desktops™ that are static or dedicated machines.

This behavior is the default on VM boot. The VDI is left in the state it was at the last shutdown.

Select this option if you plan to allow users to make permanent changes to their desktops. To select persist, shut down the VM, and then enter the following command:

```
1 xe vdi-param-set uuid=vdi_uuid on-boot=persist
```

Reset

Tip:

Use this boot behavior if you are hosting Citrix Virtual Desktops that are shared or randomly allocated machines.

On VM boot, the VDI is reverted to the state it was in at the previous boot. Any changes made while the VM is running are lost when the VM is next booted.

Select this option if you plan to deliver standardized desktops that users cannot permanently change. To select reset, shut down the VM, and then enter the following command:

Warning:

After you change `on-boot=reset`, any data saved to the VDI is discarded after the next shutdown/start or reboot.

Make the ISO library available to XenServer® hosts

To make an ISO library available to XenServer hosts, create an external NFS or SMB/CIFS share directory. The NFS or SMB/CIFS server must allow root access to the share. For NFS shares, allow access by

setting the `no_root_squash` flag when you create the share entry in `/etc/exports` on the NFS server.

Then either use XenCenter® to attach the ISO library, or connect to the host console and run the command:

```
1 xe-mount-iso-sr host:/volume
```

For advanced use, you can pass extra arguments to the mount command.

To make a Windows SMB/CIFS share available to the host, either use XenCenter, or connect to the host console and run the following command:

```
1 xe-mount-iso-sr unc_path -t cifs -o username=myname/myworkgroup
```

Replace back slashes in the `unc_path` argument with forward-slashes. For example:

```
1 xe-mount-iso-sr //server1/myisos -t cifs -o username=johndoe/mydomain
```

After mounting the share, any available ISOs are available from the **Install from ISO Library or DVD drive** list in XenCenter. These ISOs are also available as CD images from the CLI commands.

Attach the ISO to an appropriate Windows template.

Connect to a Windows VM by using Remote Desktop

You can use one of the following ways of viewing a Windows VM console, both of which support full use of the keyboard and mouse.

- Using XenCenter. This method provides a standard graphical console and uses the VNC technology built in to XenServer to provide remote access to your virtual machine console.
- Connecting using Windows Remote Desktop. This method uses the Remote Desktop Protocol technology

In XenCenter on the **Console** tab, there is a **Switch to Remote Desktop** button. This button disables the standard graphical console within XenCenter, and switches to using Remote Desktop.

If you do not have Remote Desktop enabled in the VM, this button is disabled. To enable it, install the XenServer VM Tools for Windows. Follow the procedure below to enable it in each VM that you want to connect using Remote Desktop.

To enable Remote Desktop on a Windows VM:

1. Open **System** by clicking the **Start** button, right-click on **Computer**, and then select **Properties**.
2. Click **Remote settings**. If you're prompted for an administrator password, type the password you created during the VM setup.

3. In the **Remote Desktop** area, click the check box labeled **Allow connections from computers running any version of Remote Desktop**.
4. To select any non-administrator users that can connect to this Windows VM, click the **Select Remote Users** button and provide the user names. Users with Administrator privileges on the Windows domain can connect by default.

You can now connect to this VM using Remote Desktop. For more information, see the Microsoft Knowledge Base article, [Connect to another computer using Remote Desktop Connection](#).

Note:

You cannot connect to a VM that is asleep or hibernating. Set the settings for sleep and hibernation on the remote computer to **Never**.

Time handling in Windows VMs

For Windows guests, initially the control domain clock drives the time. The time updates during VM lifecycle operations such as suspend and reboot. We recommend configuring a reliable NTP service for the control domain and all Windows VMs to use. For more information, see [Change the NTP configuration for a server](#).

If you manually set a VM to be two hours ahead of the control domain, then it persists. You might set the VM ahead by using a time-zone offset within the VM. If you later change the control domain time (either manually or by NTP), the VM shifts accordingly but maintains the two hours offset. Changing the control domain time-zone does not affect VM time-zones or offset. XenServer uses the hardware clock setting of the VM to synchronize the VM. XenServer does not use the system clock setting of the VM.

When performing suspend and resume operations or using live migration, ensure that you have up-to-date XenServer VM Tools for Windows installed. XenServer VM Tools for Windows notify the Windows kernel that a time synchronization is required after resuming (potentially on a different physical host).

Note:

If you are running Windows VMs in Citrix Virtual Desktops environment, you must ensure that the host clock has the same source as the Active Directory (AD) domain. Failure to synchronize the clocks can cause the VMs to display an incorrect time and cause the Windows PV drivers to crash.

Time handling in Linux VMs

In addition to the behavior defined by XenServer, operating system settings and behaviors can affect the time handling behavior of your Linux VMs. Some Linux operating systems might periodically

synchronize their system clock and hardware clock, or the operating system might use its own NTP service by default. For more information, see the documentation for the operating system of your Linux VM.

Note:

When installing a new Linux VM, ensure that you change the time-zone from the default UTC to your local value. For specific distribution instructions, see [Linux Release Notes](#).

Hardware clocks in Linux VMs are **not** synchronized to the clock running on the control domain and can be altered. When the VM first starts, the control domain time is used to set the initial time of the hardware clock and system clock.

If you change the time on the hardware clock, this change is persisted when the VM reboots.

System clock behavior depends on the operating system of the VM. For more information, see the documentation for your VM operating system.

You cannot change this XenServer time handling behavior.

Install VMs from Reseller Option Kit (BIOS-locked) media

There are two types of VM: BIOS-generic and BIOS-customized. To enable installation of Reseller Option Kit (BIOS-locked) OEM versions of Windows onto a VM, copy the BIOS strings of the VM from the host with which the media was supplied. Alternatively, advanced users can set user-defined values to the BIOS strings.

BIOS-generic

The VM has generic XenServer BIOS strings.

Note:

If a VM doesn't have BIOS strings set when it starts, the standard XenServer BIOS strings are inserted into it and the VM becomes BIOS-generic.

BIOS-customized

You can customize the BIOS in two ways: Copy-Host BIOS strings and User-Defined BIOS strings.

Note:

After you first start a VM, you cannot change its BIOS strings. Ensure that the BIOS strings are correct before starting the VM for the first time.

Copy-Host BIOS strings The VM has a copy of the BIOS strings of a particular host in the pool. To install the BIOS-locked media that came with your host, follow the procedures given below.

Using XenCenter:

1. Click the **Copy host BIOS strings to VM** check box in the New VM Wizard.

Using the CLI:

1. Run the `vm-install copy-bios-strings-from` command. Specify the `host-uuid` as the host from which the strings are copied (that is, the host that the media was supplied with):

```
1 xe vm-install copy-bios-strings-from=host uuid \  
2   template=template name sr-name=label=name of sr \  
3   new-name-label=name for new VM
```

This command returns the UUID of the newly created VM.

For example:

```
1 xe vm-install copy-bios-strings-from=46dd2d13-5aee-40b8-ae2c-95786  
ef4 \  
2   template="win7sp1" sr-name=label=Local\ storage \  
3   new-name-label=newcentos  
4   7cd98710-bf56-2045-48b7-e4ae219799db
```

2. If the relevant BIOS strings from the host have been successfully copied into the VM, the command `vm-is-bios-customized` confirms this success:

```
1 xe vm-is-bios-customized uuid=VM uuid
```

For example:

```
1 xe vm-is-bios-customized uuid=7cd98710-bf56-2045-48b7-e4ae219799db  
2   This VM is BIOS-customized.
```

Note:

When you start the VM, it is started on the physical host from which you copied the BIOS strings.

Warning:

It is your responsibility to comply with any EULAs governing the use of any BIOS-locked operating systems that you install.

User-defined BIOS strings The user has option to set custom values in selected BIOS strings using CLI/API. To install the media in a VM with customized BIOS, follow the procedure given below.

Using the CLI:

1. Run the `vm-install` command (without `copy-bios-strings-from`):

```
1 xe vm-install template=template name sr-name-label=name of sr \  
2   new-name-label=name for new VM
```

This command returns the UUID of the newly created VM.

For example:

```
1 xe vm-install template="win7sp1" sr-name-label=Local\ storage \  
2   new-name-label=newcentos \  
3   7cd98710-bf56-2045-48b7-e4ae219799db
```

2. To set user-defined BIOS strings, run the following command before starting the VM for the first time:

```
1 xe vm-param-set uuid=VM_UUID bios-strings:bios-vendor=VALUE \  
2   bios-strings:bios-version=VALUE bios-strings:system-  
3   manufacturer=VALUE \  
4   bios-strings:system-product-name=VALUE bios-strings:system-  
5   version=VALUE \  
6   bios-strings:system-serial-number=VALUE bios-strings:enclosure  
7   -asset-tag=VALUE
```

For example:

```
1 xe vm-param-set uuid=7cd98710-bf56-2045-48b7-e4ae219799db \  
2   bios-strings:bios-vendor="vendor name" \  
3   bios-strings:bios-version=2.4 \  
4   bios-strings:system-manufacturer="manufacturer name" \  
5   bios-strings:system-product-name=guest1 \  
6   bios-strings:system-version=1.0 \  
7   bios-strings:system-serial-number="serial number" \  
8   bios-strings:enclosure-asset-tag=abk58hr
```

Notes:

- Once the user-defined BIOS strings are set in a single CLI/API call, they cannot be modified.
- You can decide on the number of parameters you want to provide to set the user-defined BIOS strings.

Warning:

It is your responsibility to:

- Comply with any EULAs and standards for the values being set in VM's BIOS.
- Ensure that the values you provide for the parameters are working parameters. Providing

incorrect parameters can lead to boot/media installation failure.

Assign a GPU to a Windows VM (for use with Citrix Virtual Desktops)

XenServer enables you to assign a physical GPU in the XenServer host to a Windows VM running on the same host. This GPU pass-through feature benefits graphics power users, such as CAD designers, who require high performance graphics capabilities. It is supported only for use with Citrix Virtual Desktops.

While XenServer supports only one GPU for each VM, it automatically detects and groups identical physical GPUs across hosts in the same pool. Once assigned to a group of GPUs, a VM may be started on any host in the pool that has an available GPU in the group. When attached to a GPU, a VM has certain features that are no longer available, including live migration, VM snapshots with memory, and suspend/resume.

Assigning a GPU to a VM in a pool does not interfere with the operation of other VMs in the pool. However, VMs with GPUs attached are considered non-agile. If VMs with GPUs attached are members of a pool with high availability enabled, both features overlook these VMs. The VMs cannot be migrated automatically.

GPU pass-through can be enabled using XenCenter or the xe CLI.

Requirements

GPU pass-through is supported for specific machines and GPUs. In all cases, the IOMMU chipset feature (known as VT-d for Intel models) must be available and enabled on the XenServer host. Before enabling the GPU pass-through feature, visit the [Hardware Compatibility List](#).

Before assigning a GPU to a VM

Before you assign a GPU to a VM, put the appropriate physical GPUs in your XenServer host and then restart the machine. Upon restart, XenServer automatically detects any physical GPUs. To view all physical GPUs across hosts in the pool, use the `xe pgpu-list` command.

Ensure that the IOMMU chipset feature is enabled on the host. To do so, enter the following:

```
1 xe host-param-get uuid=uuid_of_host param-name=chipset-info param-key=iommu
```

If the value printed is **false**, IOMMU is not enabled, and GPU pass-through is not available using the specified XenServer host.

To assign a GPU to a Windows VM by using XenCenter:

1. Shut down the VM that you want to assign a GPU.
2. Open the VM properties: right-click the VM and select **Properties**.
3. Assign a GPU to the VM: Select GPU from the list of VM properties, and then select a GPU type. Click **OK**.
4. Start the VM.

To assign a GPU to a Windows VM by using the xe CLI:

1. Shut down the VM that you want to assign a GPU group by using the `xe vm-shutdown` command.
2. Find the UUID of the GPU group by entering the following:

```
1 xe gpu-group-list
```

This command prints all GPU groups in the pool. Note the UUID of the appropriate GPU group.

3. Attach the VM to a GPU group by entering the following:

```
1 xe vgpu-create gpu-group-uuid=uuid_of_gpu_group vm-uuid=uuid_of_vm
```

To ensure that the GPU group has been attached, run the `xe vgpu-list` command.

4. Start the VM by using the `xe vm-start` command.
5. Once the VM starts, install the graphics card drivers on the VM.

Installing the drivers is essential, as the VM has direct access to the hardware on the host. Drivers are provided by your hardware vendor.

Note:

If you try to start a VM with GPU pass-through on the host without an available GPU in the appropriate GPU group, XenServer prints an error.

To detach a Windows VM from a GPU by using XenCenter:

1. Shut down the VM.
2. Open the VM properties: right-click the VM and select **Properties**.
3. Detach the GPU from the VM: Select **GPU** from the list of VM properties, and then select **None** as the GPU type. Click **OK**.
4. Start the VM.

To detach a Windows VM from a GPU by using the xe CLI:

1. Shut down the VM by using the `xe vm-shutdown` command.

2. Find the UUID of the vGPU attached to the VM by entering the following:

```
1 xe vgpu-list vm-uuid=uuid_of_vm
```

3. Detach the GPU from the VM by entering the following:

```
1 xe vgpu-destroy uuid=uuid_of_vgpu
```

4. Start the VM by using the `xe vm-start` command.

Create ISO images

XenServer can use ISO images as installation media and data sources for Windows or Linux VMs. This section describes how to make ISO images from CD/DVD media.

To create an ISO on a Linux system:

1. Put the CD- or DVD-ROM disk into the drive. Ensure that the disk is not mounted. To check, run the command:

```
1 mount
```

If the disk is mounted, unmount the disk. See your operating system documentation for assistance if necessary.

2. As root, run the command

```
1 dd if=/dev/cdrom of=/path/cdimg_filename.iso
```

This command takes some time. When the operation is completed successfully, you see something like:

```
1 1187972+0 records in
2 1187972+0 records out
```

Your ISO file is ready.

To create an ISO on a Windows system:

Windows computers do not have an equivalent operating system command to create an ISO. Most CD-burning tools have a means of saving a CD as an ISO file.

Troubleshoot VM problems

June 24, 2026

If you see unusual VM behavior, this section aims to help you solve the problem. This section describes where application logs are located and other information that can help your XenServer® Solution Provider track and resolve the issue.

Important:

Follow the troubleshooting information in this section only under the guidance of your XenServer Solution Provider or the Support Team.

Vendor Updates: Keep your VMs up-to-date with operating system vendor-supplied updates. The vendor might have provided fixes for VM crashed and other failures.

VM crashes

If you are experiencing VM crashes, it is possible that a kernel crash dump can help identify the problem. Reproduce the crash, if possible, and follow this procedure. Consult your guest OS vendor for further investigation on this issue.

The crashdump behavior of your VMs can be controlled by using the `actions-after-crash` parameter. The following are the possible values:

Value	Description
<code>preserve</code>	Leave the VM in a paused state. (For analysis)
<code>restart</code>	No core dump, reboot VM. (This is the default)
<code>destroy</code>	No core dump, leave VM halted.

To enable saving of VM crash dumps:

1. On the XenServer host, determine the UUID of the desired VM by running the following command:

```
1 xe vm-list name-label=<name> params=uuid --minimal
```

2. Change the `actions-after-crash` value by using `xe vm-param-set`; for example, run the following command on dom0:

```
1 xe vm-param-set uuid=<vm_uuid> actions-after-crash=preserve
```

3. Crash the VM.

- a) Determine the VM's domain ID by running the following command on dom0:

```
1 xe vm-param-get uuid=<vm_uuid> param-name=dom-id
```

- b) Run the `xl trigger` command in dom0 to trigger the crash:

```
1 xl trigger <dom_id> nmi
```

Windows VM crashdump behavior

By default Windows crash dumps are put into %SystemRoot%\Minidump in the Windows VM itself. You can configure the VMs dump level by following the menu path **My Computer > Properties > Advanced > Startup and Recovery**.

Gather XenServer VM Tools for Windows diagnostics

The XenServer VM Tools for Windows include a diagnostic tool called `xtbugtool`. This tool is available as a PowerShell script, `xtbugtool.ps1`.

In addition, we provide a deprecated batch script version for running on older versions of Windows, `xtbugtool.bat`. Where possible, we recommend that you use the PowerShell version.

This tool is supported on the following operating systems:

- Windows 10
- Windows 11
- Windows Server 2016
- Windows Server 2019
- Windows Server 2022
- Windows Server 2025

To help diagnose any issues with XenServer VM Tools for Windows, the script gathers the following information from your Windows VM:

- Registry Keys relevant to XenServer VM Tools for Windows installations
- MSinfo in NFO and text versions
- System, Security and Application event logs
- The following files and directories:
 - `C:\ProgramData\Citrix`
 - `C:\ProgramData\Citrix Systems, Inc.`
 - `C:\ProgramData\XenServer`
 - `C:\Windows\Inf\setupapi.dev.log`
 - `C:\Windows\Inf\setupapi.setup.log`
- The output of the following commands:
 - `systeminfo`

- tasklist /NH
- pnputil -e
- pnputil /enum-drivers
- winmgmt /verifyrepository
- ipconfig /all

Running the PowerShell script

To run the PowerShell script, follow these steps:

1. On your Windows VM, open a PowerShell terminal as Administrator.
2. At the **Start** menu type `PowerShell` and select, **Run as Administrator**.
3. Change directory to `C:\Program Files\XenServer\XenTools\diagnostics`
4. Run the script with the following command: `xtbugtool.ps1 <path_to_output_directory>`

If you do not provide a value for the output path, the script outputs to the current directory.

The script creates a compressed file with a name in the format `xtbugtool-yyyymmddhhmmss.zip` in the output directory that you specify.

Running the batch script (deprecated)

To run the batch script, follow these steps:

1. On your Windows VM, open a command prompt as Administrator.
2. At the **Start** menu type `cmd` and select, **Run as Administrator**.
3. Change directory to `C:\Program Files\XenServer\XenTools\diagnostics`
4. Run the script with the following command:

```
xtbugtool.bat <path_to_output_directory>
```

The script creates a compressed file with a name in the format `xt-bugtool-YYYY.MM.DD-xxxx.zip` in the output directory that you specify.

UEFI and Secure Boot problems

How do I change the screen resolution of the XenCenter® console on a UEFI-enabled VM?

To change the screen resolution of the XenCenter console on a UEFI-enabled VM:

1. Open the **Windows Settings**
2. Click **Update & Security** (Windows 10/Server 2019/2022) or **System > Recovery** (Windows 11/Server 2025).
3. Under the recovery tab, press the **Restart now** button.
4. Navigate to **Troubleshoot > Advanced Options > UEFI firmware settings**.
5. Press **Restart**. During restart, the UEFI settings menu loads.
6. Navigate to **Device Manager > OVMF Platform Configuration**. This displays the current screen resolution.
7. Press **Enter** to see the screen resolution options.
8. Use the arrow keys to select the desired screen resolution and press **Enter**.
9. Press **F10** to save the changes and confirm your choice.
10. Reboot the VM to see the XenCenter console in an updated screen resolution.

Why can't I create a UEFI Secure Boot VM?

Check that your VM operating system supports UEFI Secure Boot mode. The following operating systems support Secure Boot:

- Windows 10
- Windows 11
- Windows Server 2016
- Windows Server 2019
- Windows Server 2022
- Windows Server 2025
- Red Hat Enterprise Linux 8
- Red Hat Enterprise Linux 9
- Red Hat Enterprise Linux 10 (preview)
- Ubuntu 22.04
- Ubuntu 24.04
- Rocky Linux 8
- Rocky Linux 9
- SUSE Linux Enterprise 15
- Debian Bookworm 12
- Oracle Linux 8

Why is my UEFI Secure Boot VM failing to start?

If you see the following messages on the console of your UEFI Secure Boot VM and an alert in XenCenter, the Secure Boot process has failed and your VM does not start.

```

UEFI Interactive Shell v2.2
EDK II
UEFI v2.70 (EDK II, 0x00010000)
Mapping table
  FS0: Alias(s) :F1::BLK3:
        PciRoot (0x0) /Pci (0x3,0x0) /VenHw (3D3CA290-B9A5-11E3-B75D-B8AC6F7D65E6,0
1004016) /VenMedia (C5BD4D42-1A76-4996-B956-73CDA326CD0A)
  BLK0: Alias(s) :
        PciRoot (0x0) /Pci (0x3,0x0) /VenHw (3D3CA290-B9A5-11E3-B75D-B8AC6F7D65E6,0
1000003)
  BLK1: Alias(s) :
        PciRoot (0x0) /Pci (0x3,0x0) /VenHw (3D3CA290-B9A5-11E3-B75D-B8AC6F7D65E6,0
1004016)
  BLK2: Alias(s) :
        PciRoot (0x0) /Pci (0x3,0x0) /VenHw (3D3CA290-B9A5-11E3-B75D-B8AC6F7D65E6,0
1004016) /CDROM (0x0)
Press ESC in 1 seconds to skip startup.nsh or any other key to continue.
Shell> _

```

This is usually caused by the installation of unsigned drivers into the VM. Investigate what drivers have been updated or installed since the last successful Secure Boot.

You can [disable Secure Boot](#) to remove the unsigned drivers.

Important:

Before doing this, back up your VM by taking a snapshot.

Is Secure Boot causing an issue on my VM?

To diagnose whether an issue on your VM is caused by Secure Boot being enabled for the VM, [disable Secure Boot](#) and try to reproduce the issue.

How do I run Windows debug on a Secure Boot Windows VM?

You cannot run Windows debug on a Secure Boot Windows VM. To run Windows debug on your VM, you will need to [disable Secure Boot](#).

Why are only two NICs showing up for my UEFI-enabled Windows VM?

Even if you set up more than two NICs when you created your UEFI-enabled VM, when the VM first starts you only see two NICs. After the XenServer VM Tools for Windows have been installed in the VM, this information displays correctly.

Why are my emulated devices showing as different types than expected on a UEFI Windows VM?

UEFI Secure Boot VMs use NVME and E1000 for emulated devices. However, when the VM first starts the emulated devices show as different types. After the XenServer VM Tools for Windows have been installed in the VM, this information displays correctly.

Why can't I convert my templates from BIOS mode to UEFI or UEFI Secure Boot mode?

You can only create a UEFI-enabled VM template from a template supplied with XenServer.

Do not use the `xe template-param-set` command for templates that have something installed on them or templates that you created from a snapshot. The boot mode of these snapshots cannot be changed and, if you attempt to change the boot mode, the VM fails to boot.

How do I check UEFI and UEFI Secure Boot variables?

On the XenServer host where the UEFI or UEFI Secure Boot VM is hosted, run the following commands:

```
1 varstore-ls
```

This command lists the GUIDs and names of the available variables. Use the GUID and name in the following command:

```
1 varstore-get <VM\_ID> <GUID> <name> | hexdump -C
```

Why can't I use a 'test' driver with a Secure Boot VM?

If you are also working with a third party to debug and fix issues in their UEFI Secure Boot VM, the third party provide might provide unsigned drivers for test or verification purpose. These drivers do not work in a UEFI Secure Boot VM.

Request a signed driver from the third party. Or you can switch your UEFI Secure Boot VM into setup mode to run with the unsigned driver.

Xentop utility

The xentop utility displays real-time information about a XenServer system and running domains in a semi-graphical format. You can use this tool to investigate the state of the domain associated with a VM.

To run the xentop utility:

1. Connect to the XenServer host over SSH or, in XenCenter, go to the **Console** tab of the host.
2. Run the following command: `xentop`

The console displays information about the host in a table. The information is periodically re-freshed.

Output columns

The xentop utility displays the following columns in the console:

- **NAME** - The name of the domain. “Domain-0” is the XenServer control domain. Other domains belong to the VMs.
- **STATE** - The state of the domain. The state can have one of the following values:
 - d - the domain is dying
 - s - the domain is shutting down
 - b - the domain is blocked
 - c - the domain has crashed
 - p - the domain is paused
 - r - the domain is actively running on one of the CPUs
- **CPU(sec)** - The CPU usage of the domain in seconds
- **CPU(%)** - The CPU usage of the domain as a percentage
- **MEM(k)** - The current memory usage of the domain in KiB
- **MEM(%)** - The current memory usage of the domain as a percentage
- **MAXMEM(k)** - The maximum domain memory usage in KiB
- **MAXMEM(%)** - The maximum domain memory usage as a percentage
- **VCPUS** - The number of virtual CPUs assigned to the domain
- **NETS** - The number of virtual networks used by the domain
- **NETTX(k)** - The amount of total network tx in KiB
- **NETRX(k)** - The amount of total network rx in KiB
- **VBDS** - The number of virtual block devices
- **VBD_OO** - The total number of times that the VBD has encountered an out of requests error. When that occurs, I/O requests for the VBD are delayed.
- **VBD_RD** - The total number of VBD read requests

- **VBD_WR** - The total number of VBD write requests
- **VBD_RSECT** - The VBD read sectors
- **VBD_WSECT** - The VBD write sectors

Xentop parameters

You can use the following parameters to configure the output for the xentop command:

- **-h** - Output the command help for the xentop command.
- **-V** - Output the version of the xentop command.
- **-d** or **–delay=SECONDS** - Set the number of seconds between updates
- **-n** or **–networks** - Output the data for each VIF network associated with a domain
- **-x** or **–vbds** - Output the data for each VBD block device associated with a domain
- **-r** or **repeat-header** - Repeat the table header before each domain
- **-v** or **–vcpus** - Output the data for each vCPU associated with a domain
- **-i** or **–iterations** - Number of iterations (updates) to display before xentop exits
- **-f** or **–full-name** - Output the full domain name instead of a truncated name

You can also configure most of these parameters from within the xentop utility.

High availability

March 16, 2026

The XenServer® high availability (HA) feature ensures that your VMs keep running with minimal downtime and no data corruption. When issues such as host hardware failures and network disruption occur, the XenServer pool responds by restarting affected VMs on stable hosts in the pool. This feature enables you to keep your VMs running until you are able to rectify any hardware issues.

XenServer ensures VM high availability by the following means:

- Using a network heartbeat to verify connectivity between hosts in the pool.
- Using a storage heartbeat to verify connectivity between the hosts and the shared storage.
- Detecting whether a host has failed.
- Detecting whether a host or hosts has become unreachable.
- Fencing hosts that cannot communicate with the largest partition of hosts in the pool. A fenced host restarts immediately, causing all VMs running on it to be stopped. When it has restarted it tries to rejoin the resource pool. This precaution prevents VMs from running on two hosts at once and risking data corruption.

- Marking a failed, fenced, or unreachable host as no longer part of the live set of hosts in the pool.
- Marking any VMs that were running on that host as halted.
- If the failed, fenced, or unreachable host is the pool coordinator, reassigning the coordinator role to another host in the pool.
- Restarting any halted VMs according to the failover plan you configured.
- Monitoring any changes to the pool configuration to check that the configured failover plan can be enacted.

Because the HA feature automatically restarts VMs on other hosts in the pool, XenServer must be sure that the original (failed or unreachable) host of the VM is no longer running the VM. Two instances of the same VM running at the same time can cause VM data corruption. To guard against this possibility, XenServer hosts in an HA-enabled pool are proactive about self-fencing if they are in situations that might lead to two instances of the same VM running.

The HA feature keeps your key VMs running until you can resolve the underlying hardware or network issue. When you realize a HA event has occurred, investigate and resolve the underlying failure to return your pool to full capacity.

This article describes high availability concepts, requirements, and expected behaviors. For information about configuring and managing high availability, see [Configure high availability](#).

Requirements

To use the high availability feature, you need the following items in your environment:

- **A XenServer pool:** The HA feature operates within a single resource pool.
 - We recommend that the pool is homogenous. Each host in the pool exposes the same set of CPU features to the VMs and makes it easier for the VMs to be restarted anywhere in the pool.
 - For the heartbeat mechanism to work effectively, we recommend that the pool has at least 3 hosts.
 - Ensure that all hosts in the pool are online before enabling HA.
- **Shared storage for all hosts in the pool:** To enable any VM in the pool to be restarted on any host in the pool after a failure, all hosts in the pool must have access to the SR where the VM disks are stored.
- **A heartbeat SR:** This SR can be the same SR as the one where the VM disks are stored. The pool stores information that enables the pool to coordinate failure detection and recovery in the event of a failure.
 - The heartbeat SR must be on an iSCSI, NFS, or Fibre Channel LUN. Storage attached using either SMB or iSCSI when authenticated using CHAP cannot be used as the heartbeat SR.

We recommend that this SR is highly reliable with low latency.

- XenServer 8.4 requires 4 GB for the heartbeat SR.

The information stored on the heartbeat SR includes:

- 4 MB heartbeat volume: Provides the storage heartbeat, which verifies that hosts in the pool have access to the storage.
- The metadata volume: Stores the pool coordinator metadata to be used if there is a failover of the pool coordinator. This volume takes up the rest of the required space.
- **Reliable and redundant storage communication for the heartbeat SR:** For the HA feature to have the most accurate view of which hosts can access the shared storage, configure your environment to ensure that storage traffic is reliable. For iSCSI and Fibre Channel SRs, configure multipathing. For NFS SRs, use a resilient bonded network as your storage network.
- **Static IP addresses for all hosts:** HA treats a change of host IP address as the host losing connection and assumes that the host's network has failed. As a result, the host can fence. Avoid this by only using static IPs in your pool.
- **A dedicated bonded interface on the management network:** For the HA feature to have the most accurate view of the pool state, you require reliable and redundant network communications between hosts.
- **The management network allows network heartbeat UDP traffic over port 694:** The network heartbeat verifies that the hosts in the pool are live and can contact each other.

To protect a VM running in your high availability pool, set up your VM with the following configuration:

- Store the VM disks on shared storage available to all hosts in the pool.
- Set up its virtual network interfaces on pool-wide networks.
- Ensure that the VM can use live migration. For more information, see [Migration compatibility requirements](#).
- Do not connect the VM to a local DVD drive.

A VM that fulfills all these criteria is called agile.

A VM that uses NVIDIA vGPU or GPU pass-through cannot be protected by HA. However, the HA mechanism can attempt to restart this VM on a best-effort basis.

Requirements for GFS2 clustered pools

The high availability behavior for GFS2 clustered pools uses a different underlying mechanism and as such has some different requirements and behaviors. For more information, see [GFS2 clustered pools](#).

HA failover plan

The HA mechanism calculates a pool-wide failover plan based on the following criteria:

- **VM recovery requirements:** Each VM can have a restart priority and start order defined.
- **Available pool resources:** The main resource that is considered is host memory.
- **The number of host failures to tolerate:** After you enable HA in your pool, XenServer can compute the maximum number of hosts that can fail in the pool before protected VMs cannot be restarted. You can set the number of host failures to tolerate to less than or equal to this value.

If a failover plan meeting these criteria cannot be calculated, the pool is considered overcommitted. If the protected VMs cannot be restarted in the pool, XenServer raises a system alert. This alert is also shown in the XenCenter® **Notifications** panel.

For every VM in your pool, you can define its recovery behavior.

Restart priority

You can assign a VM one of the following restart priorities:

- **Protected:** If the VM or its host goes unexpectedly offline, HA restarts the VM on another host. This restart is guaranteed, provided the pool isn't overcommitted and the VM is agile. If the VM restart fails, HA attempts to start the VM when there is extra capacity in the pool. This value is `restart` on the xe CLI and **Restart** in XenCenter.
- **Best-effort:** If the host running the VM goes unexpectedly offline, HA attempts to restart the VM on another host. It makes this attempt only after all protected VMs have been successfully restarted. High availability makes only one attempt to restart a best-effort VM. If this attempt fails, high availability does not make further attempts to restart the VM. This value is `best-effort` on the xe CLI and **Restart if possible** in XenCenter.
- **Unprotected:** If the VM or its host goes unexpectedly offline, HA does not attempt to restart the VM. This is the default setting. This value is an empty string on the xe CLI and **Do not restart** in XenCenter.

High availability never stops or migrates a running VM to free resources to restart a VM with a higher restart priority.

Start order

The start order is the order in which XenServer high availability attempts to restart protected VMs when a failure occurs. This value is used for protected VMs only. The default value is 0, which is the highest priority. Protected VMs with a start order value of 0 are restarted first. The higher the start order value, the later in the sequence the VM is restarted.

Pool behavior

After you have enabled HA in your XenServer pool, the pool exhibits the following behaviors.

Behavior during setup

When you enable HA in a pool, the pool coordinator performs the following setup:

- Calculates the initial failover plan.
- Configures the database to write updates to the heartbeat SR. This setting ensures that VM configuration changes are not lost when a host fails.
- Sets up the pool coordinator metadata on the heartbeat SR.

All pool members:

- Send network heartbeats to each other. As a result, there is a small increase in management network traffic as hosts in the pool verify that they can communicate with each other. This network traffic continues while HA is enabled.

Behavior during normal operation

During normal operation, the pool coordinator of an HA pool performs the following actions (in addition to its usual functions):

- Dynamically maintains a failover plan. This plan details what to do when a set of hosts in a pool fail at any given time. This plan takes into account the maximum number of host failures that can be tolerated and ensures that all protected VMs can be restarted. The plan is dynamically recalculated based on VM lifecycle operations and movement. If changes (for example, the addition of new VMs to the pool) mean that all protected VMs can no longer be restarted after the maximum number of host failures, a plan cannot be calculated and the pool is overcommitted. When the pool becomes overcommitted, XenServer raises an alert via XenCenter, email, SNMP trap, or NRPE alert.

During normal operation, each member of an HA pool performs the following actions (in addition to their usual functions):

- Checks that the pool coordinator is alive. The host does this by attempting to acquire a “master lock” on the shared storage. If a pool coordinator already exists, this attempt fails
- Send a network heartbeat. This network heartbeat is sent using UDP over port 694 on the management network to all other hosts in the pool.
- Maintains a record of the liveness of hosts in the pool. The liveness of hosts according to each individual host is the set of other hosts it believes to be live. If a host has not received a network

heartbeat from another host within the period specified by the HA timeout (by default, 60 seconds), it communicates with the other hosts in the pool to agree whether the liveset must be updated.

- Writes to the state file on the storage heartbeat volume. This action verifies that the host still has access to the storage. It also allows the hosts to communicate their state to one another (in addition to the communication with the network heartbeat).
- Updates the database on the heartbeat SR. The host records any changes to VM configurations for the VMs they are hosting.

Some pool operations are blocked or not advised when HA is enabled. Temporarily disable high availability to perform these operations:

- Adding a host to the pool.
- Removing a host from the pool. Blocked if this action can cause the pool to become overcommitted.
- Shutting down a host in the pool. Blocked if this action can cause the pool to become overcommitted.
- Changing the management network.
- Changing the SR attached to the pool.
- Enabling clustering. Some high availability behavior and requirements are different for clustered pools. For more information, see [Clustered pools](#).

During normal operation, performing these actions in the pool does not activate the HA failover plan:

- VM clean shutdown from XenCenter or the xe CLI. The HA mechanism does not consider this VM to have failed and does not attempt to restart it. For more information about this action, see [Shut down a VM protected by high availability](#)
- VM crash or clean shutdown from within the guest OS, if HA has been configured to not automatically reboot VMs on internal shutdown. In this case, the HA mechanism does not consider the VM to have failed and does not attempt to restart it. For more information about this setting, see [Configure the restart behavior for VMs that are shut down internally](#)
- Host clean shutdown from XenCenter or the xe CLI. The HA mechanism does not consider this host to have failed and does not attempt to restart any VMs that were hosted on it. However, if this action causes the pool to become overcommitted, it is blocked by XenServer. For more information about this action, see [Shut down a host when high availability is enabled](#)

Behavior during hardware failures or infrastructure instability

During this phase, all hosts in the pool are responsible for detecting their own connectivity status and agreeing the connectivity status of other hosts in the pool.

XenServer HA detects and handles the following types of failure:

- **Failed host or hosts:** In this situation, all remaining hosts notice very quickly that the failed host or hosts have stopped updating the state file and are no longer sending network heartbeats. After an appropriate delay, these hosts are removed from the liveset.
- **Network partition:** In this situation, one or more hosts cannot communicate with one or more other hosts. A host notices that it has not received network heartbeats from one or more other hosts within the defined timeout and starts a fault handler. This fault handler process communicates through the state file and working network heartbeat, and uses that information to determine what network partitions (groups of hosts which can talk to each other) exist. The hosts in the largest partition are the liveset and survive. If there are equal sized partitions, the hosts in the partition containing the host with the lowest host UUID survive.
- **Failed storage connection:** In this situation, a host notices that it cannot reach the storage or other hosts notice that its updates are not present on the storage. The hosts communicate through the network heartbeat communications to check whether other hosts have lost storage access:
 - If all hosts have lost storage, but not network, this is considered to be a temporary loss of storage and the hosts remain up to wait for the storage to come back. Any further failures and all hosts in the pool fence. This rule prevents storage being a single point of failure.
 - If only some hosts have lost storage access, but all hosts still have network access, these hosts are removed from the liveset.

If a host knows that it will appear failed or unreachable to the majority of the pool, that host self-fences. Fencing is an expected behavior designed as a protective measure for VM data. It ensures that a VM is not running in two places at once. A host uses the following criteria to decide that it needs to self-fence:

- If the host's toolstack is not running and cannot be restarted, the host self-fences.
- If the host has lost both the network and storage heartbeats, the host considers itself unreachable and self-fences.
- If the host has lost the storage heartbeat, but is still receiving network heartbeats:
 - If the host can still contact all other pool members and all those members have also lost the storage heartbeat, the host stays alive. This case prevents the storage acting as a single point of failure and fencing the whole pool.
 - If the host can't contact one or more other hosts in the pool, it self-fences.
- If the host has lost any network heartbeats, but still has the storage heartbeat, it determines whether it is in the largest network partition. If it is not, the host self-fences.
- There is the chance that a network communication failure might split the pool into partitions of equal size. If, using the information in the state file on the heartbeat SR, a host knows that it is in such a network partition:

- If the partition contains the host with the lowest UUID, the host stays alive.
- If the partition does not contain the host with the lowest UUID, the host self-fences.

When a fence action is taken, the host restarts immediately and abruptly, causing all VMs running on it to be stopped. The fenced host enters a reboot sequence, and when it has restarted it tries to rejoin the resource pool.

Behavior during recovery

If the pool coordinator is the host that has failed, fenced, or become unreachable, other hosts attempt to get the master lock. The host that succeeds becomes the new coordinator.

Hosts that have self-fenced restart and attempt to rejoin the pool.

When a host is marked as dead and its VMs halted, the pool coordinator is responsible for the following recovery actions.

- Restart all protected VMs according to the failover plan.
- If there is not enough resource to start all protected VMs, the pool coordinator waits until resource becomes available (for example, if previously fenced hosts rejoin the pool) and then attempts to start the protected VMs.
- After all protected VMs are successfully started, the pool coordinator makes one attempt to restart each best-effort VM.

Configure high availability

December 17, 2025

High availability attempts to protect your VM workload in the case of host or hardware failure.

This article describes the tasks to perform to configure high availability. For information about high availability concepts, requirements, and expected behaviors, see [High availability](#).

1. Configure your pool for high availability

For a pool to be compatible with high availability, it must fulfill certain requirements:

1. Ensure that your pool is homogenous, contains three or more hosts, and all the hosts are online.
For more information, see [Create a pool](#).

2. Set up a dedicated, bonded interface for the pool management network and ensure that this network allows UDP traffic over port 694. For more information, see [Create NIC bonds in resource pools](#).
3. Configure static IP addresses for all hosts in the pool.
4. Set up shared storage for all hosts in the pool. For more information, see [Create an SR](#).
5. Ensure that one shared storage repository in the pool is on an iSCSI, NFS, or Fibre Channel LUN that meets the following requirements:
 - 4 GB or more in size.
 - Has resilient storage communication: For iSCSI and Fibre Channel SRs, configure multipathing. For NFS SRs, use a resilient bonded network as your storage network.

This SR is used as the heartbeat SR.

2. Prepare your VMs for high availability

Complete these steps for any VMs in your pool that you want to protect with high availability:

1. Ensure that the VM is agile:
 - a) Ensure that its virtual disks are on shared storage.

If the disks are not on shared storage, use the following command to move them:

```
xe vm-copy uuid=<vm_uuid> new-name-label=<name_for_copy> sr-uuid=<uuid_of_sr>
```

Set the `sr-uuid` to the UUID of the pool shared storage.
 - b) Ensure that the VM is not connected to a local DVD drive.
 - i. Identify whether the VM does not have empty CD/DVD drives by typing the following:

```
1 xe vbd-list type=CD empty=false vm-uuid=<vm_uuid>
```

If there is an attached drive, this returns its information. Note the `uuid`, which is the first item in the list.
 - ii. To empty the CD/DVD drives of the VM, type the following:

```
1 xe vbd-eject uuid=<uuid>
```
 - c) Ensure that its virtual network interfaces are on pool-wide networks.
 - d) Ensure that the VM meets the requirements for live migration. For more information, see [Migration requirements](#).

2. Specify a set of priorities that determine which VMs are given the highest restart priority when a pool is overcommitted.

Set the value of `ha-restart-priority` to one of the following options:

- `restart` to have the VM restarted on another host in the pool, provided the pool isn't overcommitted. XenServer retries this restart until it succeeds.
- `best-effort` to make a single attempt to restart the VM on another host in the pool. XenServer makes this attempt only after all protected VMs have been successfully restarted.

When setting the value for `order` of the VMs, set the highest priority VMs to have the lowest start order. The default value for this setting is 0 and VMs with this start order are started first.

```
1 xe vm-param-set uuid=<vm_uuid> ha-restart-priority=<priority>
  order=<start_order>
```

You can instead use XenCenter to configure these settings on your VMs. For more information, see [Start options](#).

3. Enable high availability on your XenServer® pool

You can enable high availability on a pool by using either XenCenter or the command-line interface (CLI). For information about using XenCenter to enable high availability, see [Enable high availability](#).

To enable high availability by using the xe CLI:

1. To enable high availability on the pool and, optionally, specify a timeout, run the following command:

```
1 xe pool-ha-enable heartbeat-sr-uuids=<sr uuid> ha-config:timeout=<
  timeout_in_seconds>
```

The timeout configured by this command only applies to this enablement of high availability. If you don't specify a time-out, the default is 60 seconds. To change this default timeout for your pool, see [Configure the high availability timeout](#).

2. Compute the maximum number of hosts that can fail before there are insufficient resources to run all the protected VMs in the pool by running the following command:

```
1 xe pool-ha-compute-max-host-failures-to-tolerate
```

The number returned is how many host failures are currently possible in the pool without loss of the liveness guarantee for protected VMs. This value can change as conditions in the pool change.

3. Specify the number of host failures to tolerate that is less than or equal to the value given by the previous step.

```
1 xe pool-param-set ha-host-failures-to-tolerate=<failure_value>
   uuid=<pool uuid>
```

The number of failures to tolerate determines when an alert is sent. The system recomputes a failover plan as the state of the pool changes. It uses this computation to identify the pool capacity and how many more failures are possible without loss of the liveness guarantee for protected VMs. A system alert is generated when this computed value falls below the specified value.

Configure the high availability timeout

The timeout is the period during which networking or storage is not accessible by the hosts in your pool. If any XenServer host is unable to access networking or storage within the timeout period, it can self-fence and restart. The default timeout is 60 seconds. However, you can change this value by setting a default high availability timeout for your pool:

```
1 xe pool-param-set uuid=<pool uuid> other-config:default_ha_timeout=<
   timeout in seconds>
```

If you enable high availability by using XenCenter® instead of the xe CLI, this default still applies.

Remove high availability protection from a VM

To disable high availability features for a VM, use the following command:

```
1 xe vm-param-set uuid=<vm_uuid> ha-restart-priority=
```

This command retains the start order settings. You can enable high availability for this VM again by setting the `ha-restart-priority` parameter to `restart` or `best-effort` as appropriate.

Configure the restart behavior for VMs that are shut down internally

If a VM is protected by HA, it is restarted by default whenever it shuts down. This behavior includes internal shutdowns due to either a VM crash or a clean shutdown from within the guest operating system.

You can change this behavior to ensure that VMs in the HA-enabled pool are only restarted if the host they are running on fails or is fenced, and that in cases of internal shutdown a VM remains shut down.

To configure HA-protected VMs to not automatically reboot after an internal shutdown, set the `ha-reboot-vm-on-internal-shutdown` value to **false**:

```
1 xe pool-param-set uuid=<pool_uuid> ha-reboot-vm-on-internal-shutdown=  
  false
```

To configure HA-protected VMs to automatically reboot after an internal shutdown, set the `ha-reboot-vm-on-internal-shutdown` value to **true**:

```
1 xe pool-param-set uuid=<pool_uuid> ha-reboot-vm-on-internal-shutdown=  
  true
```

Recover an unreachable host

If for some reason, a host cannot access the high availability state file, it is possible that a host might become unreachable. To recover your XenServer installation, you might have to disable high availability using the `host-emergency-ha-disable` command on the host:

```
1 xe host-emergency-ha-disable --force
```

If the host was the pool coordinator, it starts up as normal with high availability disabled. Pool members reconnect and automatically disable high availability. If the host was a pool member and cannot contact the pool coordinator, you might have to take one of the following actions:

- Force the host to reboot as a pool coordinator (`xe pool-emergency-transition-to-master`)

```
1 xe pool-emergency-transition-to-master uuid=<host uuid>
```

- Tell the host where the new pool coordinator is (`xe pool-emergency-reset-master`):

```
1 xe pool-emergency-reset-master master-address=<new pool  
  coordinator hostname>
```

When all hosts have successfully restarted, re-enable high availability:

```
1 xe pool-ha-enable heartbeat-sr-uuid=<sr uuid>
```

Shut down a host when high availability is enabled

Take special care when shutting down or rebooting a host to prevent the high availability mechanism from assuming that the host has failed. To shut down a host cleanly when high availability is enabled, disable the host, evacuate the host, and finally shutdown the host by using either XenCenter or the CLI. To shut down a host in an environment where high availability is enabled, run these commands:

```
1 xe host-disable host=<host name>
2 xe host-evacuate uuid=<host uuid>
3 xe host-shutdown host=<host name>
```

Shut down a VM protected by high availability

When a VM is protected under a high availability plan and set to restart automatically, it cannot be shut down while this protection is active. To shut down a VM, first disable its high availability protection and then run the CLI command.

```
1 xe vm-param-set uuid=<vm_uuid> ha-restart-priority=
2 xe vm-shutdown uuid=<vm_uuid>
```

XenCenter offers you a dialog box to automate disabling the protection when you select the **Shut-down** button of a protected VM.

Note:

If you shut down a VM from within the guest, and the VM is protected, it is automatically restarted under the high availability failure conditions. The automatic restart helps ensure that an operator error doesn't result in a protected VM being left shut down accidentally. If you want to shut down this VM, disable its high availability protection first.

Disaster recovery and backup

December 17, 2025

The XenServer Disaster Recovery (DR) feature allows you to recover virtual machines (VMs) and vApps from a failure of hardware which destroys a whole pool or site. For protection against single host failures, see [High availability](#).

Note:

You must be logged on with your *root* account or have the role of *Pool Operator* or higher to use the DR feature.

Understanding XenServer® DR

XenServer DR works by storing all the information required to recover your business-critical VMs and vApps on storage repositories (SRs). The SRs are then replicated from your primary (production) environment to a backup environment. When a protected pool at your primary site goes down, you can

recover the VMs and vApps in that pool from the replicated storage recreated on a secondary (DR) site with minimal application or user downtime.

The **Disaster Recovery** settings in XenCenter® can be used to query the storage and import selected VMs and vApps to a recovery pool during a disaster. When the VMs are running in the recovery pool, the recovery pool metadata is also replicated. The replication of the pool metadata allows any changes in VM settings to be populated back to the primary pool when the primary pool recovers. Sometimes, information for the same VM can be in several places. For example, storage from the primary site, storage from the disaster recovery site and also in the pool that the data is to be imported to. If XenCenter finds that the VM information is present in two or more places, it ensures that it uses only the most recent information.

The Disaster Recovery feature can be used with XenCenter and the xe CLI. For CLI commands, see [Disaster recovery commands](#).

Tip:

You can also use the Disaster Recovery settings to run test failovers for non-disruptive testing of your disaster recovery system. In a test failover, all the steps are the same as failover. However, the VMs and vApps are not started up after they have been recovered to the disaster recovery site. When the test is complete, cleanup is performed to delete all VMs, vApps, and storage recreated on the DR site.

XenServer VMs consist of two components:

- Virtual disks that are being used by the VM, stored on configured storage repositories (SRs) in the pool where the VMs are located.
- Metadata describing the VM environment. This information is required to recreate the VM if the original VM is unavailable or corrupted. Most metadata configuration data is written when the VM is created and is updated only when you change the VM configuration. For VMs in a pool, a copy of this metadata is stored on every host in the pool.

In a DR environment, VMs are recreated on a secondary site using the pool metadata and configuration information about all VMs and vApps in the pool. The metadata for each VM includes its name, description and Universal Unique Identifier (UUID), and its memory, virtual CPU, and networking and storage configuration. It also includes VM startup options –start order, delay interval, high availability, and restart priority. The VM startup options are used when restarting the VM in a high availability or DR environment. For example, when recovering VMs during disaster recovery, VMs within a vApp are restarted in the DR pool in the order specified in the VM metadata, and using the specified delay intervals.

DR infrastructure requirements

Set up the appropriate DR infrastructure at both the primary and secondary sites to use XenServer DR.

- Storage used for pool metadata *and* the virtual disks used by the VMs must be replicated from the primary (production) environment to a backup environment. Storage replication such as using mirroring varies between devices. Therefore, consult your storage solution vendor to handle Storage replication.
- After the VMs and vApps that you recovered to a pool on your DR site are up and running, the SRs containing the DR pool metadata and virtual disks must be replicated. Replication allows the recovered VMs and vApps to be restored back to the primary site (*failed back*) when the primary site is back online.
- The hardware infrastructure at your DR site does not have to match the primary site. However, the XenServer environment must be at the same release and patch level.
- The hosts and pools at the secondary site must have the same license edition as those at the primary site. These XenServer licenses are in addition to those assigned to hosts at the primary site.
- Sufficient resources must be configured in the target pool to allow all the failed over VMs to be recreated and started.

Warning:

The Disaster Recovery settings do not control any Storage Array functionality.

Users of the Disaster Recovery feature must ensure that the metadata storage is, in some way replicated between the two sites. Some Storage Arrays contain “Mirroring” features to achieve the replication automatically. If you use these features, you must disable the mirror functionality (“mirror is broken”) before restarting VMs on the recovery site.

Deployment considerations

Review the following steps before enabling Disaster Recovery.

Steps to take before a disaster

The following section describes the steps to take before disaster.

- Configure your VMs and vApps.

- Note how your VMs and vApps are mapped to SRs, and the SRs to LUNs. Take particular care with the naming of the `name_label` and `name_description` parameters. Recovering VMs and vApps from replicated storage is easier if the names of SRs capture how VMs and vApps are mapped to SRs, and SRs to LUNs.
- Arrange replication of the LUNs.
- Enable pool metadata replication to one or more SRs on these LUNs.
- Ensure that the SRs you are replicating the primary pool metadata to are attached to only one pool.

Steps to take after a disaster

The following section describes the steps to take after a disaster has occurred.

- Break any existing storage mirrors so that the recovery site has read/write access to the shared storage.
- Ensure that the LUNs you want to recover VM data from are not attached to any other pool, or corruption can occur.
- If you want to protect the *recovery* site from a disaster, you must enable pool metadata replication to one or more SRs on the recovery site.

Steps to take after a recovery

The following section describes the steps to take after a successful recovery of data.

- Resynchronize any storage mirrors.
- On the recovery site, cleanly shut down the VMs or vApps that you want to move back to the primary site.
- On the primary site, follow the same procedure as for the failover in the previous section, to failback selected VMs or vApps to the primary
- To protect the primary site against future disaster - you must re-enable pool metadata replication to one or more SRs on the replicated LUNs.

Enable Disaster Recovery

March 18, 2026

This section describes how to enable Disaster Recovery in XenCenter. Use the **Configure DR** option to identify storage repositories where the pool metadata, configuration information about all the VMs and vApps in the pool is stored. The metadata is updated whenever you change the VM or vApp configuration within the pool.

Note:

You can enable Disaster Recovery only when using LVM over HBA or LVM over iSCSI. A small amount of space is required on this storage for a new LUN which contains the pool recovery information.

Before you begin, ensure that the SRs used for DR are attached only to the pool at the primary site. SRs used for DR must not be attached to the pool at the secondary site.

To configure Disaster Recovery, complete the following steps:

1. On the primary site, select the pool that you want to protect. From the **Pool** menu, point to **Disaster Recovery**, and then select **Configure**.
2. Select up to 8 SRs where the pool metadata can be stored. A small amount of space is required on this storage for a new LUN which contains the pool recovery information.

Note:

Information for all VMs in the pool is stored, VMs do not need to be independently selected for protection.

3. Select **OK**. Your pool is now protected.

Recover VMs and vApps during a disaster (Failover)

This section explains how to recover your VMs and vApps on the secondary (recovery) site.

1. In XenCenter select the secondary pool, and on the **Pool** menu, select **Disaster Recovery** and then **Disaster Recovery Wizard**.

The **Disaster Recovery** wizard displays three recovery options: **Failover**, **Failback**, and **Test Failover**. To recover on to your secondary site, select **Failover** and then select **Next**.

Warning:

If you use Fibre Channel shared storage with LUN mirroring to replicate data to the secondary site, break the mirroring before attempting to recover VMs. Mirroring must be broken to ensure that the secondary site has Read/Write access.

2. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover.

By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, select **Find Storage Repositories** and then select the storage type to scan for:

- To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
- To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then type the target host, IQN, and LUN details.

When you have selected the required SRs in the wizard, select **Next** to continue.

3. Select the VMs and vApps that you want to recover. Select the appropriate **Power state after recovery** option to specify whether you want the wizard to start them up automatically when they have been recovered. Alternatively, you can start them up manually after failover is complete.

Select **Next** to progress to the next wizard page and begin failover pre-checks.

4. The wizard performs several pre-checks before starting failover. For example, to ensure that all the storage required by the selected VMs and vApps is available. If any storage is missing at this point, you can select **Attach SR** on this page to find and attach the relevant SR.

Resolve any issues on the pre-checks page, and then select **Failover** to begin the recovery process.

5. A progress page displays the result of the recovery process for each VM and vApp. The Failover process exports the metadata for VMs and vApps from the replicated storage. Therefore, the time taken for Failover depends on the VMs and vApps you recover. The VMs and vApps are recreated in the primary pool, and the SRs containing the virtual disks are attached to the recreated VMs. If specified, the VMs are started.
6. When the failover is complete, select **Next** to see the summary report. Select **Finish** on the summary report page to close the wizard.

When the primary site is available, work through the **Disaster Recovery** wizard and select **Failback** to return to running your VMs on that site.

Restore VMs and vApps to the primary site after disaster (Failback)

This section explains how to restore VMs and vApps from replicated storage. You can restore VMs and vApps back to a pool on your primary (production) site when the primary site comes back up after a disaster. To failback VMs and vApps to your primary site, use the **Disaster Recovery** wizard.

1. In XenCenter select the primary pool, and on the Pool menu, select **Disaster Recovery** and then **Disaster Recovery Wizard**.

The **Disaster Recovery** wizard displays three recovery options: **Failover**, **Failback**, and **Test Failover**. To restore VMs and vApps to your primary site, select **Failback** and then select **Next**.

Warning:

When you use Fibre Channel shared storage with LUN mirroring to replicate data to the primary site, break the mirroring before attempting to restore VMs. Mirroring must be broken to ensure that the primary site has Read/Write access.

2. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover.

By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, choose **Find Storage Repositories** and then select the storage type to scan for:

- To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
- To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then type the target host, IQN, and LUN details.

When you have selected the required SRs in the wizard, select **Next** to continue.

3. Select the VMs and vApps that you want to restore. Select the appropriate **Power state after recovery** option to specify whether you want the wizard to start them up automatically when they have been recovered. Alternatively, you can start them up manually after failback is complete.

Select **Next** to progress to the next wizard page and begin failback pre-checks.

4. The wizard performs several pre-checks before starting failback. For example, to ensure that all the storage required by the selected VMs and vApps is available. If any storage is missing at this point, you can select **Attach SR on this page** to find and attach the relevant SR.

Resolve any issues on the pre-checks page, and then select **Failback** to begin the recovery process.

5. A progress page displays the result of the recovery process for each VM and vApp. The Failback process exports the metadata for VMs and vApps from the replicated storage. Therefore, Failback can take some time depending on the number of VMs and vApps you are restoring. The VMs and vApps are recreated in the primary pool, and the SRs containing the virtual disks are attached to the recreated VMs. If specified, the VMs are started.
6. When the failback is complete, select **Next** to see the summary report. Select **Finish** on the summary report page to close the wizard.

Test failover

Failover testing is an essential component in disaster recovery planning. You can use the **Disaster Recovery** wizard to perform non-disruptive testing of your disaster recovery system. During a test failover operation, the steps are the same as for failover. However, instead of being started after they have been recovered to the DR site, the VMs and vApps are placed in a paused state. At the end of a test failover operation, all VMs, vApps, and storage recreated on the DR site are automatically deleted. After initial DR configuration, and after you make significant configuration changes in a DR-enabled pool, verify that failover works correctly by performing a test failover.

1. In XenCenter select the secondary pool, and on the **Pool** menu, select **Disaster Recovery** to open the **Disaster Recovery Wizard**.

The **Disaster Recovery** wizard displays three recovery options: **Failover**, **Failback**, and **Test Failover**. To test your disaster recovery system, select **Test Failover** and then select **Next**.

Note:

If you use Fibre Channel shared storage with LUN mirroring to replicate data to the secondary site, break the mirroring before attempting to recover data. Mirroring must be broken to ensure that the secondary site has Read/Write access.

2. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover.

By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, select **Find Storage Repositories** and then the storage type to scan for:

- To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
- To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then type the target host, IQN, and LUN details in the box.

When you have selected the required SRs in the wizard, select **Next** to continue.

3. Select the VMs and vApps that you want to recover then select **Next** to progress to the next page and begin failover pre-checks.
4. Before beginning the test failover, the wizard performs several pre-checks. For example, to ensure that all the storage required by the selected VMs and vApps is available.
 - **Check that storage is available.** If any storage is missing, you can select Attach SR on this page to find and attach the relevant SR.
 - **Check that high availability is not enabled on the target DR pool.** High availability must be disabled on the secondary pool to avoid having the same VMs running on both the primary and DR pools. High availability must be disabled to ensure that it does not start the

recovered VMs and vApps automatically after recovery. To disable high availability on the secondary pool, you can simply select **Disable HA** on the page. If high availability is disabled at this point, it is enabled again automatically at the end of the test failover process.

Resolve any issues on the pre-checks page, and then select **Failover** to begin the test failover.

5. A progress page displays the result of the recovery process for each VM and vApp. The Failover process recovers metadata for the VMs and vApps from the replicated storage. Therefore, Failover can take some time depending on the number of VMs and vApps you are recovering. The VMs and vApps are recreated in the DR pool, the SRs containing the virtual disks are attached to the recreated VMs.

The recovered VMs are placed in a paused state: they do not start up on the secondary site during a test failover.

6. After you are satisfied that the test failover was performed successfully, select Next in the wizard to have the wizard clean up on the DR site:
 - VMs and vApps that were recovered during the test failover are deleted.
 - Storage that was recovered during the test failover is detached.
 - If high availability on the DR pool was disabled at the pre-checks stage to allow the test failover to take place, it is re-enabled automatically.

The progress of the cleanup process appears on the wizard.

7. Select **Finish** to close the wizard.

vApps

December 17, 2025

A vApp is logical group of one or more related Virtual Machines (VMs). vApps can be started up as a single entity when there is a disaster. When a vApp is started, the VMs contained within the vApp start in a user predefined order. The start order allows VMs which depend upon one another to be automatically sequenced. An administrator no longer has to manually sequence the startup of dependent VMs when a whole service requires restarting. For example, during a software update. The VMs within the vApp do not have to reside on one host and are distributed within a pool using the normal rules. The vApp feature is useful in the Disaster Recovery (DR) situation. In a DR scenario, an Administrator may group all VMs on the same Storage Repository, or which relate to the same Service Level Agreement (SLA).

To group VMs together in a vApp follow the procedure:

1. Select the pool and, on the **Pool** menu, click **Manage vApps**.

2. Type a name for the vApp, and optionally a description, and then click **Next**.

You can choose any name you like, but an informative name is best. Although we recommend that you avoid having multiple vApps using the same name, it is not a requirement. XenCenter® does not enforce any constraints regarding unique vApp names. It is not necessary to use quotation marks for names that include spaces.

3. Select which VMs to include in the new vApp, and then click **Next**.

You can use the search option to list only VMs with names that include the specified text string.

4. Specify the startup sequence for the VMs in the vApp, and then click **Next**.

Start Order: Specifies the order in which individual VMs are started within the vApp, allowing certain VMs to be restarted before others. VMs with a start order value of 0 (zero) are started first. VMs with a start order value of 1 are started next, and then the VMs with a value of 2, and so on.

Attempt to start next VM after: A delay interval that specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence.

5. You can review the vApp configuration on the final page. Click **Previous** to go back and change any settings, or **Finish** to create the vApp.

Note:

A vApp can span multiple hosts in a single pool, but cannot span across several pools.

Manage vApps in XenCenter

The **Manage vApps** setting in XenCenter allows you to create, delete, and change vApps. It also enables you to start and shut down vApps, and import and export vApps within the selected pool. When you select a vApp in the list, the VMs it contains are listed in the details pane. For more information, see [vApps](#) in the XenCenter documentation.

Back up and restore hosts and VMs

December 17, 2025

Whenever possible, leave the installed state of XenServer hosts unaltered. That is, do not install any additional packages or start additional services on XenServer hosts and treat them as appliances. The best way to restore, then, is to reinstall XenServer host software from the installation media. If you

have multiple XenServer hosts, the best approach is to configure a TFTP server and appropriate answer files for this purpose. For more information, see [Network boot installations](#).

We recommend that you use a backup solution offered by a third-party provider. For more information, see [Third-party backup solutions](#).

XenServer Premium Edition customers can take advantage of the faster changed block only backup. For more information, see the [changed block tracking documentation](#).

These solutions are certified by their vendors and not by Cloud Software Group. Support for these backup solutions is provided by the solution vendor.

We recommend that you frequently perform as many of the following backup procedures as possible to recover from possible server and software failure.

To back up pool metadata:

1. Run the command:

```
1 xe pool-dump-database file-name=backup
```

2. To restore the database, run the command:

```
1 xe pool-restore-database file-name=backup dry-run=true
```

This command checks that the target machine has an appropriate number of appropriately named NICs, which is required for the backup to succeed.

To back up host configuration and software:

1. Run the command:

```
1 xe host-backup host=host file-name=hostbackup
```

Notes:

- Do not create the backup in the control domain.
- The backup procedure can create a large backup file.
- To complete a restore, you must reboot to the original install CD.
- This data can only be restored to the original machine.

To back up a VM:

1. Ensure that the VM to be backed up is offline.
2. Run the command:

```
1 xe vm-export vm=vm_uuid filename=backup
```

Note:

This backup also backs up all of the VM data. When importing a VM, you can specify the storage mechanism to use for the backed-up data.

Warning:

The backup process can take longer to complete as it backs up all of the VM data.

To back up VM metadata only:

Run the command:

```
1 xe vm-export vm=vm_uuid filename=backup metadata=true
```

Back up virtual machine metadata

XenServer® hosts use a database on each host to store metadata about VMs and associated resources such as storage and networking. When combined with SRs, this database forms the complete view of all VMs available across the pool. Therefore it is important to understand how to back up this database to recover from physical hardware failure and other disaster scenarios.

This section first describes how to back up metadata for single-host installations, and then for more complex pool setups.

Back up single host installations

Use the CLI to back up the pool database. To obtain a consistent pool metadata backup file, run `pool-dump-database` on the XenServer host and archive the resulting file. The backup file contains sensitive authentication information about the pool, so ensure it is securely stored.

To restore the pool database, use the `xe pool-restore-database` command from a previous dump file. If your XenServer host has died completely, then you must first do a fresh install, and then run the `pool-restore-database` command against the freshly installed XenServer host.

After you restore the pool database, some VMs may still be registered as being **Suspended**. However, if the storage repository with the suspended memory state defined in the `suspend-VDI-uuid` field, is a local SR, then the SR may not be available as the host has been reinstalled. To reset these VMs back to the **Halted** state so that they can start up again, use the `xe vm-shutdown vm=vm_name -force` command, or use the `xe vm-reset-powerstate vm=vm_name -force` command.

Warning:

XenServer preserves UUIDs of the hosts restored using this method. If you restore to a different physical machine while the original XenServer host is still running, duplicate UUIDs may be

present. As a result, XenCenter refuses to connect to the second XenServer host. Pool database backup is not the recommended mechanism for cloning physical hosts. Use the automated installation support instead. For more information, see [Install](#).

Back up pooled installations

In a pool scenario, the pool coordinator provides an authoritative database that is synchronously mirrored to all the pool member hosts. This process provides a level of built-in redundancy to a pool. Any pool member can replace the pool coordinator because each pool member has an accurate version of the pool database. For more information on how to transition a member into becoming a pool coordinator, see [Hosts and resource pools](#).

This level of protection may not be sufficient. For example, when shared storage containing the VM data is backed up in multiple sites, but the local server storage (containing the pool metadata) is not. To re-create a pool given a set of shared storage, you must first back up the `pool-dump-database` file on the pool coordinator host, and archive this file. To restore this backup later on a brand new set of hosts:

1. Install a fresh set of XenServer hosts from the installation media, or if applicable, network boot from your TFTP server.
2. Use the `xe pool-restore-database` on the host designated to be the new pool coordinator.
3. Run the `xe host-forget` command on the new pool coordinator to remove the old member machines.
4. Use the `xe pool-join` command on the member hosts to connect them to the new pool.

Back up XenServer hosts

This section describes the XenServer host control domain backup and restore procedures. These procedures do *not* back up the storage repositories that house the VMs, but only the privileged control domain that runs Xen® and the XenServer agent.

Note:

The privileged control domain is best left as installed, without customizing it with other packages. We recommend that you set up a network boot environment to install XenServer cleanly from the XenServer media as a recovery strategy. Typically, you do not need to back up the control domain, but we recommend that you save the pool metadata (see [Back up virtual machine metadata](#)). Consider this backup method as complementary to backing up the pool metadata.

Using the `xe` commands `host-backup` and `host-restore` is another approach that you can take. The `xe host-backup` command archives the active partition to a file you specify. The `xe host-restore` command extracts an archive created by `xe host-backup` over the currently inactive disk partition of the host. This partition can then be made active by booting off the installation CD and selecting to restore the appropriate backup.

After completing the steps in the previous section and rebooting the host, ensure that the VM metadata is restored to a consistent state. Run `xe pool-restore-database` on `/var/backup/pool-database-{ DATE }` to restore the VM metadata. This file is created by `xe host-backup` using `xe pool-dump-database` command before archiving the running filesystem, to snapshot a consistent state of the VM metadata.

To back up your XenServer host:

On a remote host with enough disk space, run the following command

```
1 xe host-backup file-name=filename -h hostname -u root -pw password
```

This command creates a compressed image of the control domain file system. The image is stored in the location specified by the `file-name` argument.

To restore a running XenServer host:

1. If you want to restore your XenServer host from a specific backup, run the following command while the XenServer host is up and reachable:

```
1 xe host-restore file-name=filename -h hostname -u root -pw
  password
```

This command restores the compressed image back to the hard disk of the XenServer host which runs this command (not the host on which `filename` resides). In this context, “restore” may be a misnomer, as the word usually suggests that the backed-up state has been put fully in place. The restore command only unpacks the compressed backup file and restores it to its normal form. However, it is written to another partition (`/dev/sda2`) and does *not* overwrite the current version of the filesystem.

2. To use the restored version of the root filesystem, reboot the XenServer host using the XenServer installation CD and select the **Restore from backup** option.

After the restore from backup is completed, reboot the XenServer host and it will start up from the restored image.

3. Finally, restore the VM metadata using the following command:

```
1 xe pool-restore-database file-name=/var/backup/pool-database-* -h
  hostname -u root -pw password
```

Note:

Restoring from a backup as described in this section does not destroy the backup partition.

To restart a crashed XenServer host:

If your XenServer host has crashed and is not reachable, use the XenServer installation CD to do an upgrade install. When the upgrade install is complete, reboot the machine and ensure that your host is reachable with XenCenter® or remote CLI.

Then proceed with backing up XenServer hosts as describes in this section.

Back up VMs

We recommend that you use a backup solution offered by a third-party provider. For more information, see [Third-party backup solutions](#).

These solutions are certified by their vendors and not by Cloud Software Group. Support for these backup solutions is provided by the solution vendor.

XenServer Premium Edition customers can take advantage of the faster changed block only backup. For more information, see [Changed Block Tracking](#).

VM snapshots

December 17, 2025

XenServer provides a convenient mechanism that can take a snapshot of a VM storage and metadata at a given time. Where necessary, I/O is temporarily halted while the snapshot is being taken to ensure that a self-consistent disk image can be captured.

Snapshot operations result in a snapshot VM that is similar to a template. The VM snapshot contains all the storage information and VM configuration, including attached VIFs, allowing them to be exported and restored for backup purposes. Snapshots are supported on all storage types. However, for the LVM-based storage types the following requirements must be met:

- If the storage repository was created on a previous version of XenServer®, it must have been upgraded
- The volume must be in the default format (you cannot take a snapshot of `type=raw` volumes)

The snapshot operation is a two-step process:

- Capturing metadata as a template.

- Creating a VDI snapshot of the disks.

The following types of VM snapshots are supported: regular and snapshot with memory.

Regular snapshots

Regular snapshots are crash consistent and can be performed on all VM types, including Linux VMs.

Snapshots with memory

In addition to saving the VMs memory (storage) and metadata, snapshots with memory also save the VMs state (RAM). This feature can be useful when you upgrade or patch software, but you also want the option to revert to the pre-change VM state (RAM). Reverting to a snapshot with memory, does not require a reboot of the VM.

You can take a snapshot with memory of a running or suspended VM by using the management API, the xe CLI, or XenCenter.

Create a VM snapshot

Before taking a snapshot, see the following information about any special operating system-specific configuration and considerations:

- [Prepare to clone a Windows VM by using Sysprep](#)
- [Prepare to clone a Linux VM](#)

First, ensure that the VM is running or suspended so that the memory status can be captured. The simplest way to select the VM on which the operation is to be performed is by supplying the argument `vm=name` or `vm=vm uuid`.

Run the `vm-snapshot` command to take a snapshot of a VM.

```
1 xe vm-snapshot vm=vm uuid new-name-label=vm_snapshot_name
```

Create a snapshot with memory

Run the `vm-checkpoint` command, giving a descriptive name for the snapshot with memory, so that you can identify it later:

```
1 xe vm-checkpoint vm=vm uuid new-name-label=name of the checkpoint
```

When XenServer has completed creating the snapshot with memory, its UUID is displayed.

For example:

```
1 xe vm-checkpoint vm=2d1d9a08-e479-2f0a-69e7-24a0e062dd35 \  
2 new-name-label=example_checkpoint_1 \  
3 b3c0f369-59a1-dd16-ecd4-a1211df29886
```

A snapshot with memory requires at least 4 MB of disk space per disk, plus the size of the RAM, plus around 20% overhead. So a checkpoint with 256 MB RAM would require approximately 300 MB of storage.

Note:

During the checkpoint creation process, the VM is paused for a brief period, and cannot be used during this period.

To list all of the snapshots on your XenServer pool

Run the `snapshot-list` command:

```
1 xe snapshot-list
```

This command lists all of the snapshots in the XenServer pool.

To list the snapshots on a particular VM

Get the UUID of the particular VM by running the `vm-list` command.

```
1 xe vm-list
```

This command displays a list of all VMs and their UUIDs. For example:

```
1 xe vm-list  
2 uuid ( RO): 116dd310-a0ef-a830-37c8-df41521ff72d  
3 name-label ( RW): Windows Server 2016 (1)  
4 power-state ( RO): halted  
5  
6 uuid ( RO): dff45c56-426a-4450-a094-d3bba0a2ba3f  
7 name-label ( RW): Control domain on host  
8 power-state ( RO): running
```

VMs can also be specified by filtering the full list of VMs on the values of fields.

For example, specifying `power-state=halted` selects all VMs whose power-state field is equal to 'halted'. Where multiple VMs are matching, the option `--multiple` must be specified to perform the operation. Obtain the full list of fields that can be matched by using the command `xe vm-list params=all`.

Locate the required VM and then enter the following:

```
1 xe snapshot-list snapshot-of=vm uuid
```

For example:

```
1 xe snapshot-list snapshot-of=2d1d9a08-e479-2f0a-69e7-24a0e062dd35
```

This command lists the snapshots currently on that VM:

```
1      uuid ( RO): d7eefb03-39bc-80f8-8d73-2ca1bab7dcff
2      name-label ( RW): Regular
3      name-description ( RW):
4      snapshot_of ( RO): 2d1d9a08-e479-2f0a-69e7-24a0e062dd35
5      snapshot_time ( RO): 20090914T15:37:00Z
6
7      uuid ( RO): 1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef
8      name-label ( RW): Snapshot with memory
9      name-description ( RW):
10     snapshot_of ( RO): 2d1d9a08-e479-2f0a-69e7-24a0e062dd35
11     snapshot_time ( RO): 20090914T15:39:45Z
```

Restore a VM to its previous state

Ensure that you have the UUID of the snapshot that you want to revert to, and then run the `snapshot -revert` command:

1. Run the `snapshot-list` command to find the UUID of the snapshot or checkpoint that you want to revert to:

```
1 xe snapshot-list
```

2. Note the UUID of the snapshot, and then run the following command to revert:

```
1 xe snapshot-revert snapshot-uuid=snapshot uuid
```

For example:

```
1 xe snapshot-revert snapshot-uuid=b3c0f369-59a1-dd16-ecd4-
  a1211df29886
```

After reverting a VM to a checkpoint, the VM is suspended.

Notes:

- If there's insufficient disk space available to thickly provision the snapshot, you cannot restore to the snapshot until the current disk's state has been freed. If this issue occurs, retry the operation.
- It is possible to revert to any snapshot. Existing snapshots and checkpoints are not deleted

during the revert operation.

Delete a snapshot

Ensure that you have the UUID of the checkpoint or snapshot that you want to remove, and then run the following command:

1. Run the `snapshot-list` command to find the UUID of the snapshot or checkpoint that you want to revert to:

```
1 xe snapshot-list
```

2. Note the UUID of the snapshot, and then run the `snapshot-uninstall` command to remove it:

```
1 xe snapshot-uninstall snapshot-uuid=snapshot-uuid
```

3. This command alerts you to the VM and VDI that are deleted. Type `yes` to confirm.

For example:

```
1      xe snapshot-uninstall snapshot-uuid=1760561d-a5d1-5d5e-2be5-
      d0dd99a3b1ef
2      The following items are about to be destroyed
3      VM : 1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef (Snapshot with memory)
4      VDI: 11a4aa81-3c6b-4f7d-805a-b6ea02947582 (0)
5      VDI: 43c33fe7-a768-4612-bf8c-c385e2c657ed (1)
6      VDI: 4c33c84a-a874-42db-85b5-5e29174fa9b2 (Suspend image)
7      Type 'yes' to continue
8      yes
9      All objects destroyed
```

If you only want to remove the metadata of a checkpoint or snapshot, run the following command:

```
1 xe snapshot-destroy snapshot-uuid=snapshot-uuid
```

For example:

```
1 xe snapshot-destroy snapshot-uuid=d7eefb03-39bc-80f8-8d73-2ca1bab7dcff
```

Snapshot templates

Create a template from a snapshot

You can create a VM template from a snapshot. However, its memory state is removed.

1. Use the command `snapshot-copy` and specify a `new-name-label` for the template:

```
1 xe snapshot-copy new-name-label=vm-template-name \  
2     snapshot-uuid=uuid of the snapshot
```

For example:

```
1 xe snapshot-copy new-name-label=example_template_1 \  
2     snapshot-uuid=b3c0f369-59a1-dd16-ecd4-a1211df29886
```

Note:

This command creates a template object in the SAME pool. This template exists in the XenServer database for the current pool only.

2. To verify that the template has been created, run the command `template-list`:

```
1 xe template-list
```

This command lists all of the templates on the XenServer host.

Export a snapshot to a template

When you export a VM snapshot, a complete copy of the VM (including disk images) is stored as a single file on your local machine. This file has a `.xva` file name extension.

1. Use the command `snapshot-export-to-template` to create a template file:

```
1 xe snapshot-export-to-template snapshot-uuid=snapshot-uuid \  
2     filename=template- filename
```

For example:

```
1 xe snapshot-export-to-template snapshot-uuid=b3c0f369-59a1-dd16-   
    ecd4-a1211df29886 \  
2     filename=example_template_export
```

The VM export/import feature can be used in various different ways:

- As a convenient backup facility for your VMs. An exported VM file can be used to recover an entire VM in a disaster scenario.
- As a way of quickly copying a VM, for example, a special-purpose server configuration that you use many times. You simply configure the VM the way you want it, export it, and then import it to create copies of your original VM.
- As a simple method for moving a VM to another host.

For more information about the use of templates, see [Create VMs](#) and also the [Managing VMs](#) article in the XenCenter documentation.

Scheduled snapshots

The Scheduled Snapshots feature provides a simple backup and restore utility for your critical service VMs. Regular scheduled snapshots are taken automatically and can be used to restore individual VMs. Scheduled Snapshots work by having pool-wide snapshot schedules for selected VMs in the pool. When a snapshot schedule is enabled, Snapshots of the specified VM are taken at the scheduled time each hour, day, or week. Several Scheduled Snapshots may be enabled in a pool, covering different VMs and with different schedules. A VM can be assigned to only one snapshot schedule at a time.

XenCenter provides a range of tools to help you use this feature:

- To define a Scheduled Snapshot, use the **New snapshot schedule** wizard.
- To enable, disable, edit, and delete Scheduled Snapshots for a pool, use the **VM Snapshot Schedules** dialog box.
- To edit a snapshot schedule, open its **Properties** dialog box from the **VM Snapshot Schedules** dialog box.
- To revert a VM to a scheduled snapshot, select the snapshot on the **Snapshots** tab and revert the VM to it.

For more information, see [Scheduled Snapshots](#) in the XenCenter documentation.

Cope with machine failures

December 17, 2025

This section provides details of how to recover from various failure scenarios. All failure recovery scenarios require the use of one or more of the backup types listed in [Backup](#).

Member failures

In the absence of HA, pool coordinator nodes detect the failures of members by receiving regular heartbeat messages. If no heartbeat has been received for 600 seconds, the pool coordinator assumes the member is dead. There are two ways to recover from this problem:

- Repair the dead host (for example, by physically rebooting it). When the connection to the member is restored, the pool coordinator marks the member as alive again.
- Shut down the host and instruct the pool coordinator to forget about the member node using the `xe host-forget` CLI command. Once the member has been forgotten, all the VMs which were running there are marked as offline and can be restarted on other XenServer hosts.

It is important to ensure that the XenServer® host is actually offline, otherwise VM data corruption might occur.

Do not to split your pool into multiple pools of a single host by using `xe host-forget`. This action might result in them all mapping the same shared storage and corrupting VM data.

Warning:

- If you are going to use the forgotten host as an active host again, perform a fresh installation of the XenServer software.
- Do not use `xe host-forget` command if HA is enabled on the pool. Disable HA first, then forget the host, and then re-enable HA.

When a member XenServer host fails, there might be VMs still registered in the *running* state. If you are sure that the member XenServer host is definitely down, use the `xe vm-reset-powerstate` CLI command to set the power state of the VMs to *halted*. See `vm-reset-powerstate` for more details.

Warning:

Incorrect use of this command can lead to data corruption. Only use this command if necessary.

Before you can start VMs on another XenServer host, you are also required to release the locks on VM storage. Only on host at a time can use each disk in an SR. It is key to make the disk accessible to other XenServer hosts once a host has failed. To do so, run the following script on the pool coordinator for each SR that contains disks of any affected VMs: `/opt/xensource/sm/resetvdis.py host_UUID SR_UUID master`

You need only supply the third string (“master”) if the failed host was the SR pool coordinator at the time of the crash. (The SR pool coordinator is the pool coordinator or a XenServer host using local storage.)

Warning:

Be sure that the host is down before running this command. Incorrect use of this command can lead to data corruption.

If you attempt to start a VM on another XenServer host before running the `resetvdis.py` script, then you receive the following error message: `VDI <UUID> already attached RW.`

Pool coordinator failures

Every member of a resource pool contains all the information necessary to take over the role of pool coordinator if necessary. When a pool coordinator node fails, the following sequence of events occurs:

1. If HA is enabled, another pool coordinator is elected automatically.
2. If HA is not enabled, each member waits for the pool coordinator to return.

If the pool coordinator comes back up at this point, it re-establishes communication with its members, and operation returns to normal.

If the pool coordinator is dead, choose one of the members and run the command `xe pool-emergency-transition-to-master` on it. Once it has become the pool coordinator, run the command `xe pool-recover-slaves` and the members now point to the new pool coordinator.

If you repair or replace the host that was the original pool coordinator, you can simply bring it up, install the XenServer software, and add it to the pool. Since the XenServer hosts in the pool are enforced to be homogeneous, there is no real need to make the replaced host the pool coordinator.

When a member XenServer host is transitioned to being a pool coordinator, check that the default pool storage repository is set to an appropriate value. This check can be done using the `xe pool-param-list` command and verifying that the `default-SR` parameter is pointing to a valid storage repository.

Pool failures

In the unfortunate event that your entire resource pool fails, you must recreate the pool database from scratch. Be sure to regularly back up your pool-metadata using the `xe pool-dump-database` CLI command (see `pool-dump-database`).

To restore a completely failed pool:

1. Install a fresh set of hosts. Do not pool them up at this stage.
2. For the host nominated as the pool coordinator, restore the pool database from your backup using the `xe pool-restore-database` command (see `pool-restore-database`).
3. Connect to the pool coordinator by using XenCenter® and ensure that all your shared storage and VMs are available again.
4. Perform a pool join operation on the remaining freshly installed member hosts, and start up your VMs on the appropriate hosts.

Cope with failure due to configuration errors

If the physical host machine is operational but the software or host configuration is corrupted:

1. Run the following command to restore host software and configuration:

```
1 xe host-restore host=host file-name=hostbackup
```

2. Reboot to the host installation CD and select **Restore from backup**.

Physical machine failure

If the physical host machine has failed, use the appropriate procedure from the following list to recover.

Warning:

Any VMs running on a previous member (or the previous host) which have failed are still marked as **Running** in the database. This behavior is for safety. Simultaneously starting a VM on two different hosts would lead to severe disk corruption. If you are sure that the machines (and VMs) are offline you can reset the VM power state to **Halted**:

```
xe vm-reset-powerstate vm=vm_uuid --force
```

VMs can then be restarted using XenCenter or the CLI.

To replace a failed pool coordinator with a still running member:

1. Run the following commands:

```
1 xe pool-emergency-transition-to-master
2 xe pool-recover-slaves
```

2. If the commands succeed, restart the VMs.

To restore a pool with all hosts failed:

1. Run the command:

```
1 xe pool-restore-database file-name=backup
```

Warning:

This command only succeeds if the target machine has an appropriate number of appropriately named NICs.

2. If the target machine has a different view of the storage than the original machine, modify the storage configuration using the **pbid-destroy** command. Next use the **pbid-create** command to recreate storage configurations. See [pbid commands](#) for documentation of these commands.
3. If you have created a storage configuration, use **pbid-plug** or **Storage > Repair Storage Repository** menu item in XenCenter to use the new configuration.

4. Restart all VMs.

To restore a VM when VM storage is not available:

1. Run the following command:

```
1 xe vm-import filename=backup metadata=true
```

2. If the metadata import fails, run the command:

```
1 xe vm-import filename=backup metadata=true --force
```

This command attempts to restore the VM metadata on a ‘best effort’ basis.

3. Restart all VMs.

Third-party backup solutions

December 17, 2025

The following vendors have implemented support for backing up XenServer® VMs and metadata.

Support for these third-party solutions is provided by the solution vendor, not by Citrix.

If you are a vendor for a backup solution that supports XenServer and want to be added to this page, reach out to us at feedback@xenserver.com.

Vinchin

<https://www.vinchin.com/xenserver-backup.html>

This solution supports [Changed Block Tracking](#).

SEP sesam

https://wiki.sepsoftware.com/wiki/index.php/5_0_0:Citrix_XenServer_Backup

This solution supports [Changed Block Tracking](#).

Commvault

https://documentation.commvault.com/11.40/expert/backup_options_for_citrix_hypervisor_xenserver.html

XenOrchestra

<https://docs.xen-orchestra.com/backup>

This solution supports [Changed Block Tracking](#).

Veeam (BETA)

<https://forums.veeam.com/kvm-rhv-olvm-pve-schc-f62/veeam-plug-in-for-xcp-ng-12-3-2-beta2-t100399.html>

Veeam customers can download the beta and install it in non-production environments to test and provide feedback through their Veeam systems engineer. Since this is a beta release, there is no support and no upgrade path to the full version from Veeam.

This solution supports [Changed Block Tracking](#).

Hycu (BETA)

<https://www.hycu.com/blog/whats-new-hycu-r-cloud-one-global-view-stronger-threat-hunting-and-new-workload-support>

Hycu customers can download the beta and install it in non-production environments to test and provide feedback through their Hycu systems engineer. Since this is a beta release, there is no support and no upgrade path to the full version from Hycu.

This solution supports [Changed Block Tracking](#).

Workload Balancing

June 26, 2026

Notes:

- Workload Balancing is available for XenServer® Premium Edition customers. For more information about XenServer licensing, see [Licensing](#). To upgrade, or to get a XenServer license, visit the [XenServer website](#).
- Workload Balancing 8.4.1 is compatible with XenServer 8.4 and XenServer 9. If you perform a rolling pool upgrade from XenServer 8.4 to XenServer 9, update the Workload Balancing virtual appliance to ver-

sion 8.4.1 before performing the rolling pool upgrade. XenServer 9 supports only Workload Balancing 8.4.1 and later. You can download the latest version of the Workload Balancing virtual appliance from the [XenServer Downloads page](#).

Workload Balancing is a XenServer Premium Edition component, packaged as a virtual appliance, that provides the following features:

- Create reports about virtual machine (VM) performance in your XenServer environment
- Evaluate resource utilization and locate VMs on the best possible hosts in the pool for their workload's needs
- Balance VM workloads across hosts in a XenServer resource pool
- Determine the best host on which to start a VM
- Determine the best host on which to resume a VM that you powered off
- Determine the best host to move a VM to when a host fails
- Determine the optimal server for each of the host's VMs when you put a host into or take a host out of maintenance mode

Depending on your preference, Workload Balancing can accomplish these tasks automatically or prompt you to accept its rebalancing and placement recommendations. You can also configure Workload Balancing to power off hosts automatically at specific times of day. For example, configure your hosts to switch off at night to save power.

Workload Balancing can send notifications in XenCenter® regarding the actions it takes. For more information on how to configure the alert level for Workload Balancing alerts by using the xe CLI, see [Set alert level for Workload Balancing alerts in XenCenter](#).

Workload Balancing functions by evaluating the use of VMs across a pool. When a host exceeds a performance threshold, Workload Balancing relocates the VM to a less-taxed host in the pool. To rebalance workloads, Workload Balancing moves VMs to balance the resource use on hosts.

To ensure that the rebalancing and placement recommendations align with your environment's needs, you can configure Workload Balancing to optimize workloads in one of the following ways:

- To maximize resource performance
- To maximize the number of VMs that fit on hosts

These optimization modes can be configured to change automatically at predefined times or stay the same always. For

extra granularity, fine-tune the weighting of individual resource metrics: CPU, network, disk, and memory.

To help you perform capacity planning, Workload Balancing provides historical reports about host and pool health,

optimization and VM performance, and VM motion history.

As Workload Balancing captures performance data, you can also use this component to generate reports, known as

Workload Reports, about your virtualized environment. For more information, see

[Generate workload reports](#).

Workload Balancing basic concepts

When VMs are running, they consume computing resources on the physical host. These resources include CPU, Memory,

Network Reads, Network Writes, Disk Reads, and Disk Writes. Some VMs, depending on their workload, might consume more

CPU resources than other VMs on the same host. Workload is defined by the applications running on a VM and their user

transactions. The combined resource consumption of all VMs on a host reduces the available resources on the host.

Workload Balancing captures data for resource performance on VMs and physical hosts and stores it in a database.

Workload Balancing uses this data, combined with the preferences you set, to provide optimization and placement

recommendations.

Optimizations are a way in which hosts are “improved” to align with your goals: Workload Balancing makes

recommendations to redistribute the VMs across hosts in the pool to increase either performance or density. When

Workload Balancing is making recommendations, it makes them in light of its goal: to create balance or harmony across

the hosts in the pool. If Workload Balancing acts on these recommendations, the action is known as an optimization.

When Workload Balancing is enabled, XenCenter provides star ratings to indicate the optimal hosts for starting a VM.

These ratings are also provided:

- When you want to start the VM when it is powered off
- When you want to start the VM when it is suspended
- When you want to migrate the VM to a different host (Migrate and Maintenance Mode)

Within a Workload Balancing context:

- **Performance** is the usage of physical resources on a host (for example, the CPU, memory, network, and disk utilization on a host). When you set Workload Balancing to maximize performance, it recommends placing VMs to ensure that the maximum amount of resources are available for each VM.
- **Density** is the number of VMs on a host. When you set Workload Balancing to maximize density, it recommends placing VMs so you can reduce the number of hosts powered on in a pool. It ensures that the VMs have adequate computing power.

Workload Balancing does not conflict with settings you already specified for High Availability: these features are compatible.

What's new in Workload Balancing

June 26, 2026

The latest version of the Workload Balancing virtual appliance is version 8.4.1. You can download this version of the Workload Balancing virtual appliance from the [XenServer Downloads page](#).

What's new in 8.4.1

Released April 7, 2026

This update includes the following improvements:

- Removed legacy cryptographic algorithms.

- Completed end-to-end testing of Workload Balancing on XenServer® 9, including compatibility testing with new network names, regression testing of existing functionality, and security validation of the updated cipher suite.

You can use the provided migration script to migrate your data when you update to this version from Workload Balancing

8.3.0 or 8.4.0. For more information about using the migration script, see

[Migrate data from an existing virtual appliance.](#)

Get started with Workload Balancing

June 26, 2026

You can configure the Workload Balancing virtual appliance in just a few steps:

1. Review the prerequisite information and plan your Workload Balancing usage.
2. Download the Workload Balancing virtual appliance.
3. Import the Workload Balancing virtual appliance into XenCenter®.
4. Configure Workload Balancing virtual appliance from the virtual appliance console.
5. (Optional) If you already have a previous version of Workload Balancing installed, you can migrate data from an existing virtual appliance.

Note:

If you perform a rolling pool upgrade from XenServer® 8.4 to XenServer 9, update your Workload Balancing virtual appliance to version 8.4.1 before performing the rolling pool upgrade. XenServer 9 supports only Workload Balancing 8.4.1 and later. You can download the latest version of the Workload Balancing virtual appliance from the [XenServer Downloads page](#).

6. Connect your pool to the Workload Balancing virtual appliance by using XenCenter.

The Workload Balancing tab only appears in XenCenter if your pool has the required license to use Workload Balancing.

Before you start

The Workload Balancing virtual appliance is a single pre-installed VM designed to run on a XenServer host. Before

importing it, review the prerequisite information and considerations.

Prerequisites

- Workload Balancing 8.4.1 is compatible with XenServer 8.4 and XenServer 9. We recommend using the XenCenter management console to import the virtual appliance.
- If you are currently using an earlier version of the Workload Balancing virtual appliance, you can use the migration script to migrate your existing data when you upgrade to the latest version. For more information, see [Migrate from an existing virtual appliance](#).
- The Workload Balancing virtual appliance requires a minimum of 2 GB of RAM and 30 GB of disk space to run. By default, the Workload Balancing virtual appliance is assigned 2 vCPUs. This value is optimized to manage 1000 VMs within a single pool. You do not usually need to increase it. Only decrease the number of vCPUs assigned to the virtual appliance if you have a small environment. For more information, see [Change the Workload Balancing virtual appliance configuration](#).
- It is recommended to use one Workload Balancing virtual appliance to balance a single pool.

Pool requirements

To balance a pool with Workload Balancing, the pool must meet the following requirements:

- All hosts are licensed with a Premium Edition license
- All hosts meet the requirements for live migration:
 - Shared remote storage
 - Similar processor configurations
 - Gigabit Ethernet

- The pool does not contain any vGPU-enabled VMs. Workload Balancing cannot create a capacity plan for VMs that have vGPUs attached.

Considerations

Before importing the virtual appliance, note the following information and make the appropriate changes to your environment, as applicable.

- **Communications port.** Before you launch the Workload Balancing Configuration wizard, determine the port over which you want the Workload Balancing virtual appliance to communicate. You are prompted for this port during Workload Balancing Configuration. By default, the Workload Balancing server uses 8012.

Note:

Do not set the Workload Balancing port to port 443. The Workload Balancing virtual appliance cannot accept connections over port 443 (the standard TLS/HTTPS port).

- **Accounts for Workload Balancing.** There are three different accounts that are used when configuring your Workload Balancing virtual appliance and connecting it to XenServer.

The Workload Balancing Configuration wizard creates the following accounts with a user name and password that you specify:

- *Workload Balancing account*

This account is used by the XenServer host to connect to the Workload Balancing server. By default, the user name for this account is `wlbuser`. This user is created on the Workload Balancing virtual appliance during Workload Balancing configuration.

- *Database account*

This account is used to access the PostgreSQL database on the Workload Balancing virtual appliance. By default, the user name is `postgres`. You set the password for this account during Workload Balancing configuration.

When connecting the Workload Balancing virtual appliance to a XenServer pool, you must specify an existing account:

- *XenServer account*

This account is used by the Workload Balancing virtual appliance to connect to the XenServer pool and read the RRDs. Ensure that this user account has the permissions to read the XenServer pool, host, and VM RRDs. For example, provide the credentials for a user that has the `pool-admin` or `pool-operator` role.

- **Monitoring across pools.** You can put the Workload Balancing virtual appliance in one pool and monitor a different pool with it. (For example, the Workload Balancing virtual appliance is in Pool A but you are using it to monitor Pool B.)
- **Time synchronization.** The Workload Balancing virtual appliance requires that the time on the physical computer hosting the virtual appliance matches that in use by the monitored pool. There is no way to change the time on the Workload Balancing virtual appliance. We recommend pointing both the physical computer hosting Workload Balancing and the hosts in the pool it is monitoring to the same Network Time (NTP) server.
- **XenServer and Workload Balancing communicate over HTTPS.** Therefore, during Workload Balancing Configuration, Workload Balancing automatically creates a self-signed certificate on your behalf. You can change this certificate to one from a certificate authority or configure XenServer to verify the certificate or both. For information, see the [Certificates](#).
- **Storing historical data and disk space size.** The amount of historical data you can store is based on the following:
 - The size of the virtual disk allocated to Workload Balancing (by default 30 GB)
 - The minimum disk required space, which is 4,096 MB by default and controlled by the `GroomingRequiredMinimumDiskSizeInMB` parameter in the `wlb.conf` file.

The more historical data Workload Balancing collects, the more accurate and balanced the recommendations are. If

you want to store much historical data, you can do one of the following:

- Archive the data as described in [Archive database data](#)
- Make the virtual disk size assigned to the Workload Balancing virtual appliance larger as described in [Extend the virtual appliance disk](#)

For example, when you want to use the Workload Balancing Pool Audit trail feature and configure the report granularity to medium or above.

- **Load balancing Workload Balancing.** If you want to use your Workload Balancing virtual appliance to manage itself, specify shared remote storage when importing the virtual appliance.

Note:

Workload Balancing cannot perform Start On placement recommendation for the Workload Balancing virtual appliance when you are using Workload Balancing to manage itself. The reason that Workload Balancing cannot make placement recommendations when it is managing itself is because the virtual appliance must be running to perform that function. However, it can balance the Workload Balancing virtual appliance just like it would balance any other VM it is managing.

- **Plan for resource pool sizing.** Workload Balancing requires specific configurations to run successfully in large pools. For more information, see [Change the Workload Balancing virtual appliance configuration](#).

Download the virtual appliance

The Workload Balancing virtual appliance is packaged in an `.xva` format. You can download the virtual appliance from the [XenServer Downloads page](#). When downloading the file, save it to a folder on your local hard drive (typically on the computer where XenCenter is installed).

When the `.xva` download is complete, you can import it into XenCenter as described in [Import the Workload Balancing virtual appliance](#).

Import the Workload Balancing virtual appliance

Use XenCenter to import the Workload Balancing virtual appliance into a pool.

To import the virtual appliance into XenServer:

1. Open XenCenter.
2. Right-click on the pool (or host) into which you want to import the virtual appliance package, and select **Import**.
3. Browse to the `vpx-wlb.xva` package.
4. Select the pool or Home Server where you want to run the Workload Balancing virtual appliance.

When you select the pool, the VM automatically starts on the most suitable host in that pool.

Alternatively, if you don't manage the Workload Balancing virtual appliance using Workload Balancing, you can set a Home Server for the Workload Balancing virtual appliance. This setting ensures that the virtual appliance always starts on the same host.

5. Choose a storage repository on which to store the virtual disk for the Workload Balancing virtual appliance. This repository must have a minimum of 30 GB of free space.

You can choose either local or remote storage. However, if you choose local storage, you cannot manage the virtual appliance with Workload Balancing.

6. Define the virtual interfaces for the Workload Balancing virtual appliance. In this release, Workload Balancing is designed to communicate on a single virtual interface.
7. Choose a network that can access the pool you want Workload Balancing to manage.
8. Leave the **Start VMs after import** check box enabled, and click **Finish** to import the virtual appliance.
9. After you finish importing the Workload Balancing `.xva` file, the Workload Balancing VM appears in the **Resource** pane in XenCenter.

After importing the Workload Balancing virtual appliance, configure the virtual appliance as described in

Configure the Workload Balancing virtual appliance.

Configure the Workload Balancing virtual appliance

After you finish importing the Workload Balancing virtual appliance, you must configure it before you can use it to manage your pool. To guide you through the configuration, the Workload Balancing virtual appliance provides you with a configuration wizard in XenCenter. To display it, select the virtual appliance in the **Resource** pane and click the

Console tab. For all options, press **Enter** to accept the default choice.

1. After importing the Workload Balancing virtual appliance, click the **Console** tab.
2. Enter `yes` to accept the terms of the license agreement. To decline the EULA, enter `no`.

Note:

The Workload Balancing virtual appliance is also subject to the licenses contained in the `/opt/vpx/wlb` directory in the Workload Balancing virtual appliance.

3. Enter and confirm a new root password for the Workload Balancing VM. We recommend selecting a strong password.

Note:

When you enter the password, the console does not display placeholders, such as asterisks, for the characters.

4. If connecting a pool using the FQDN of the Workload Balancing virtual appliance, enter the desired computer name for the appliance. If using the IP address, leave this field blank.
5. If connecting using the FQDN, enter the domain suffix for the virtual appliance. If using the IP address, leave this field blank.

For example, if the fully qualified domain name (FQDN) for the virtual appliance is `wlb-vpx-pos-pool.domain4.bedford4.ctx`, enter `domain4.bedford4.ctx`.

Note:

The Workload Balancing virtual appliance does not automatically add its FQDN to your Domain Name System (DNS) server. Therefore, if you want the pool to use an FQDN to connect to Workload Balancing,

you must add the FQDN
to your DNS server.

6. Enter **y** to use DHCP to obtain the IP address automatically for the Workload Balancing VM. Otherwise, enter **n** and then enter a static IP address with the subnet mask as a CIDR prefix and specify the gateway for the VM.

Note:

Using DHCP is acceptable provided the lease of the IP address does not expire. It is important that the IP address does not change: When it changes, it breaks the connection between XenServer and Workload Balancing.

7. Enter a password for the Workload Balancing database user account ('postgres'). The Workload Balancing services use this account to read/write to the Workload Balancing database. Note the password that you set. You might need it later to administer to the Workload Balancing PostgreSQL database directly (for example, if you wanted to export data).
8. After pressing **Enter**, messages appear stating that the Configuration wizard is loading database objects.
9. Enter a user name and password for the Workload Balancing Server.

This action creates the account XenServer uses to connect to Workload Balancing. The default user name is **wlbuser**.
10. Enter the port for the Workload Balancing Server. The Workload Balancing server communicates by using this port.

By default, the Workload Balancing server uses 8012. The port number cannot be set to 443, which is the default TLS port number.

Note:

If you change the port here, specify that new port number when you connect the pool to Workload Balancing. For example, by specifying the port in the **Connect to WLB Server** dialog.

Ensure that the port you specify for Workload Balancing is open in any firewalls.

After you press **Enter**, Workload Balancing continues with the virtual appliance configuration, including creating self-signed certificates.

11. Now, you can also log in to the virtual appliance by entering the VM user name (typically `root`) and the root password you created earlier. However, logging in is only required when you want to run Workload Balancing commands or edit the Workload Balancing configuration file.

After configuring Workload Balancing, connect your pool to the Workload Balancing virtual appliance as described in

Connect to the Workload Balancing virtual appliance.

If necessary, you can find the Workload Balancing configuration file in the following location:

`/opt/vpx/wlb/wlb.conf`. For more information, see

[Edit the Workload Balancing configuration file](#)

The Workload Balancing log file is in this location: `/var/log/wlb/LogFile.log`. For more information, see

[Increase the detail in the Workload Balancing log](#).

Connect to the Workload Balancing virtual appliance

After configuring Workload Balancing, connect the pools you want managed to the Workload Balancing virtual appliance

by using either the CLI or XenCenter.

Note:

It is recommended to use one Workload Balancing virtual appliance to balance a single pool.

To connect a pool to your Workload Balancing virtual appliance, you need the following information:

- IP address or FQDN of the Workload Balancing virtual appliance
 - To obtain the IP address for the Workload Balancing virtual appliance:
 1. In XenCenter, go to the Workload Balancing virtual appliance **Console** tab.
 2. Log in as `root` with the root password you created when you imported the appliance.
 3. Run the following command: `ifconfig`.
 - To specify the Workload Balancing FQDN when connecting to the Workload Balancing server, first add its host name and IP address to your DNS server.

- The port number of the Workload Balancing virtual appliance. By default, XenServer connects to Workload Balancing on port 8012.

Only edit the port number when you have changed it during Workload Balancing Configuration. The port number specified during Workload Balancing Configuration, in any firewall rules, and in the **Connect to WLB Server** dialog must match.

- Credentials for the Workload Balancing account you created during Workload Balancing configuration.

This account is often known as the Workload Balancing user account. XenServer uses this account to communicate with Workload Balancing. You created this account on the Workload Balancing virtual appliance during Workload Balancing Configuration.

- Credentials for the resource pool (that is, the pool coordinator) you want Workload Balancing to monitor.

This account is used by the Workload Balancing virtual appliance to connect to the XenServer pool. This account is created on the XenServer pool coordinator and has the `pool-admin` or `pool-operator` role.

When you first connect to Workload Balancing, it uses the default thresholds and settings for balancing workloads.

Automatic features, such as Automated Optimization Mode, Power Management, and Automation, are disabled by default.

Working with certificates

If you want to upload a different (trusted) certificate or configure certificate verification, note the following

before connecting your pool to Workload Balancing:

- If you want XenServer to verify the self-signed Workload Balancing certificate, you must use the Workload Balancing IP address to connect to Workload Balancing. The self-signed certificate is issued to Workload Balancing based on its IP address.

- If you want to use a certificate from a certificate authority, it is easier to specify the FQDN when connecting to Workload Balancing. However, you can specify a static IP address in the **Connect to WLB Server** dialog. Use this IP address as the Subject Alternative Name (SAN) in the certificate.

Certificate verification is enabled by default on your pool. Disable certificate verification if your pool can't connect to Workload Balancing:

1. Get the pool UUID:

```
1 xe pool-list
```

2. Disable certificate verification:


```
1 xe pool-param-set wlb-verify-cert=false uuid=uuid_of_pool
```

If your pool can connect to Workload Balancing with certificate verification turned off, the issue is with your certificate configuration. If it cannot connect, the issue is with either your Workload Balancing credentials or your network connection. Contact Support or see [Troubleshoot Workload Balancing](#) for more information.


For more information, see [Certificates](#).

To connect your pool to the Workload Balancing virtual appliance

1. In XenCenter, select your resource pool and in its **Properties** pane, click the **WLB** tab. The **WLB** tab displays the **Connect** button.



2. In the **WLB** tab, click **Connect**. The **Connect to WLB Server** dialog box appears.



Connect to WLB Server

Server Address
Enter the address of the Workload Balancing server this Citrix Hypervisor resource pool will use.

Address:

Port: (Default is 8012)

WLB Server Credentials
Enter the credentials Citrix Hypervisor will use to connect to the Workload Balancing server.

Username:

Password:

Citrix Hypervisor Credentials
Enter the credentials the Workload Balancing Server will use to connect to Citrix Hypervisor.

Username:

Password:

☐ Use the current XenCenter credentials


OK Cancel

3. In the **Server Address** section, enter the following:

- a) In the **Address** box, type the IP address or FQDN of the Workload Balancing virtual appliance. For example,
`WLB-appliance-computername.yourdomain.net`.
- b) (Optional) If you changed the Workload Balancing port during Workload Balancing Configuration, enter the port number in the **Port** box. XenServer uses this port to communicate with Workload Balancing.

By default, XenServer connects to Workload Balancing on port 8012.

4. In the **WLB Server Credentials** section, enter the user name and password that the pool uses to connect to the Workload Balancing virtual appliance.



☒ **Update Credentials**

WLB Server Credentials
Enter the credentials Citrix Hypervisor will use to connect to the Workload Balancing server.


Username:

Password:

These credentials must be the account you created during Workload Balancing configuration. By default, the user name for this account is [wlbuser](#).

5. In the **XenServer Credentials** section, enter the user name and password for the pool you are configuring.

Workload Balancing uses these credentials to connect to the hosts in the pool.

The screenshot shows a configuration window titled "Citrix Hypervisor Credentials". Below the title is a subtitle: "Enter the credentials the Workload Balancing Server will use to connect to Citrix Hypervisor." There are two input fields: "Username:" and "Password:". To the right of these fields is an icon depicting two server racks connected by a blue arrow. Below the input fields is a checkbox labeled "Use the current XenCenter credentials".

To use the credentials with which you are currently logged into XenServer, select **Use the current XenCenter credentials**. If you have assigned a role to this account using the Access Control feature (RBAC), ensure that the role has sufficient permissions to configure Workload Balancing. For more information, see [Workload Balancing Access Control Permissions](#).

After connecting the pool to the Workload Balancing virtual appliance, Workload Balancing automatically begins monitoring the pool with the default optimization settings. To modify these settings or change the priority given to specific resources, wait at least 60 seconds before proceeding. Or wait until the XenCenter Log shows that discovery is finished.

Important:

After Workload Balancing runs for a time, if you don't receive optimal placement recommendations, evaluate your performance thresholds. This evaluation is described in [Understand when Workload Balancing makes recommendations](#). It is critical to set Workload Balancing to the correct thresholds for your environment or its recommendations might not be appropriate.

Migrate data from an existing virtual appliance

If you are using the Workload Balancing virtual appliance provided with XenServer, you can use the migration script to

migrate your existing data when you upgrade to the latest version.

The version of Workload Balancing currently provided with XenServer is 8.4.1. However, Workload Balancing 8.2.0, 8.2.1, 8.2.2, 8.3.0, and 8.4.0 were previously available. You can use this migration script to migrate from Workload Balancing 8.3.0 or 8.4.0 to Workload Balancing 8.4.1.

To use the migration script, you must have the following information:

- The root password of the existing Workload Balancing virtual appliance for remote SSH access
- The password of the data base user `postgres` on the existing Workload Balancing virtual appliance
- The password of the data base user `postgres` on the new Workload Balancing virtual appliance

Leave the existing Workload Balancing virtual appliance running on your pool while you complete the migration steps.

1. Follow the steps in the preceding section to import the new Workload Balancing virtual appliance.
2. In the SSH console of the new Workload Balancing virtual appliance, run one of the following commands.
 - For the Workload Balancing 8.3.0 virtual appliance, run:

```
1 /opt/vpx/wlb/migrate_db.sh 8.3.0 <IP of existing Workload Balancing appliance>
```

- For the Workload Balancing 8.4.0 virtual appliance, run:

```
1 /opt/vpx/wlb/migrate_db.sh 8.4.0 <IP of existing Workload Balancing appliance>
```

The command prompts you for password information when required.

3. Connect the XenServer pool with the new Workload Balancing virtual appliance.
4. After you are satisfied with the behavior of this version of the Workload Balancing virtual appliance, you can archive the old version of the virtual appliance.

Notes:

- In the case of a non-recoverable failure, reimport the latest version of the Workload Balancing virtual appliance.
- Do not disconnect the existing Workload Balancing virtual appliance. Otherwise, the data

- on the existing virtual appliance is removed.
- Keep the existing Workload Balancing virtual appliance until you have ensured that the new Workload Balancing virtual appliance is working as required.
- If necessary, you can roll back this migration by reconnecting the old version of the Workload Balancing virtual appliance to the XenServer pool.

Workload Balancing basic tasks

June 26, 2026

When you first begin using Workload Balancing, there are some basic tasks you use Workload Balancing for regularly:

- [Determining the best host on which to run a VM](#)
- [Accepting Workload Balancing optimization recommendations](#)
- [Running reports about the workloads in your environment](#)

In addition to enabling you to perform these basic tasks, Workload Balancing is a powerful XenServer® component that optimizes the workloads in your environment. The features that enable you to optimize your workloads include:

- Host power management
- Scheduling optimization-mode changes
- Running reports
- Fine-tuning the criteria Workload Balancing uses to make optimization recommendations.

For more information about these more complex features, see [Administer Workload Balancing](#).


Notes:

- Workload Balancing is available for XenServer Premium Edition customers. For more information about XenServer licensing, see [Licensing](#). To upgrade, or to get a XenServer license, visit the [XenServer website](#).
- Workload Balancing 8.4.1 is compatible with XenServer 8.4 and XenServer 9.
- The base OS of the Workload Balancing virtual appliance was originally derived from a

general-purpose Linux distribution. The Workload Balancing virtual appliance in its entirety is supported by XenServer as documented in the [XenServer Product Matrix](#).

Determine the best host on which to run a VM

When you have enabled Workload Balancing and you restart an offline VM, XenCenter® recommends the optimal pool members to start the VM on. The recommendations are also known as star ratings since stars are used to indicate the best host.



When Workload Balancing is enabled, XenCenter provides star ratings to indicate the optimal hosts for starting a VM.

These ratings are also provided:

- When you want to start the VM when it is powered off
- When you want to start the VM when it is suspended
- When you want to migrate the VM to a different host (Migrate and Maintenance Mode)

When you use these features with Workload Balancing enabled, host recommendations appear as star ratings beside the name of the physical host. Five empty stars indicate the lowest-rated, and thus least optimal, host. If you can't start or migrate a VM to a host, the host name is grayed out in the menu command for a placement feature. The reason it cannot accept the VM appears beside it.

The term *optimal* indicates the physical host best suited to hosting your workload. There are several factors

Workload Balancing uses when determining which host is optimal for a workload:

- **The amount of resources available on each host in the pool.** When a pool runs in Maximum Performance mode, Workload Balancing tries to balance the VMs across the hosts so that all VMs have good performance. When a pool runs in Maximum Density mode, Workload Balancing places VMs onto hosts as densely as possible while ensuring the VMs have sufficient resources.
- **The optimization mode in which the pool is running (Maximum Performance or Maximum Density).** When a pool runs in Maximum Performance mode, Workload Balancing places VMs on hosts with the most resources available of the type the VM requires. When a pool runs in Maximum Density mode, Workload Balancing places VMs on hosts that already have VMs running. This approach ensures that VMs run on as few hosts as possible.
- **The amount and type of resources the VM requires.** After Workload Balancing monitors a VM for a while, it uses the VM metrics to make placement recommendations according to the type of resources the VM requires. For example, Workload Balancing might select a host with less available CPU but more available memory if it is what the VM requires.

In general, Workload Balancing functions more effectively and makes better, less frequent optimization recommendations if you start VMs on the hosts it recommends. To follow the host recommendations, use one of the placement features to select the host with the most stars beside it. Placement recommendations can also be useful in Citrix Virtual Apps and Desktops™ environments.

To start a VM on the optimal host

1. In the Resources pane of XenCenter, select the VM you want to start.
2. From the VM menu, select Start on Server and then select one of the following:
 - **Optimal Server.** The optimal server is the physical host that is best suited to the resource demands of the VM you are starting. Workload Balancing determines the optimal server based on its historical records of performance metrics and your placement strategy. The optimal server is the host with the most stars.

- **One of the servers with star ratings** listed under the Optimal Server command. Five stars indicate the most-recommended (optimal) host and five empty stars indicate the least-recommended host.

Tip:

You can also select Start on Server by right-clicking the VM you want to start in the Resources pane.

To resume a VM on the optimal host

1. In the Resources pane of XenCenter, select the suspended VM you want to resume.
2. From the VM menu, select Resume on Server and then select one of the following:
 - **Optimal Server.** The optimal server is the physical host that is best suited to the resource demands of the VM you are starting. Workload Balancing determines the optimal server based on its historical records of performance metrics and your placement strategy. The optimal server is the host with the most stars.
 - **One of the servers with star ratings** listed under the Optimal Server command. Five stars indicate the most-recommended (optimal) host and five empty stars indicate the least-recommended host.

Tip:

You can also select Resume on Server by right-clicking the suspended VM in the Resources pane.

Accept Workload Balancing optimization recommendations

After Workload Balancing is running for a while, it begins to make recommendations about ways in which you can improve your environment. For example, if your goal is to improve VM density on hosts, at some point, Workload Balancing might recommend that you consolidate VMs on a host. If you aren't running in automated mode, you can choose either to accept this recommendation and apply it or to ignore it.

Important:


After Workload Balancing runs for a time, if you don't receive optimal placement recommendations, evaluate your performance thresholds. This evaluation is described in [Understand when Workload Balancing makes recommendations](#). It is critical to set Workload Balancing to the correct thresholds for your environment or its recommendations might not be appropriate.

Optimization recommendations are based on the following criteria:

- Placement strategy you select (that is, the optimization mode).
Find out the optimization mode for a pool by using XenCenter to select the pool. Look in the Configuration section of the WLB tab for the information.
- Performance metrics for resources such as a physical host's CPU, memory, network, and disk utilization.
- The role of the host in the resource pool.

When making placement recommendations, Workload Balancing considers the pool coordinator for VM placement only if no other host can accept the workload. Likewise, when a pool operates in Maximum Density mode, Workload Balancing considers the pool coordinator last when determining the order to fill hosts with VMs.

Optimization recommendations appear in the **WLB optimization** tab in XenCenter.



Optimization recommendations display the following information:

- The name of the VM that Workload Balancing recommends relocating
- The host that the VM currently resides on
- The host Workload Balancing recommends as the new location.

The optimization recommendations also display the reason Workload Balancing recommends moving the VM. For example, the recommendation displays “CPU” to improve CPU utilization. When Workload Balancing power management is enabled, Workload Balancing also displays optimization recommendations for hosts it recommends powering on or off.

Specifically, these recommendations are for consolidations.

You can click **Apply Recommendations**, to perform all operations listed in the Optimization Recommendations list.

To accept an optimization recommendation

1. In the Resources pane of XenCenter, select the resource pool for which you want to display recommendations.
2. Click the WLB tab. If there are any recommended optimizations for any VMs on the selected resource pool, they display in the Optimization Recommendations section of the WLB tab.
3. To accept the recommendations, click Apply Recommendations. XenServer begins performing all the operations listed in the Operations column of the Optimization Recommendations section.

After you click Apply Recommendations, XenCenter automatically displays the **Notifications > Events** view so you can see the progress of the VM migration.

Understand Workload Balancing recommendations under high availability

If you have Workload Balancing and XenServer High Availability enabled in the same pool, it is helpful to understand how the two features interact. Workload Balancing is designed not to interfere with High Availability. When there is a conflict between a Workload Balancing recommendation and a High Availability setting, the **High Availability** setting always takes precedence. In practice, this precedence means that:

- If attempting to start a VM on a host violates the High Availability plan, Workload Balancing doesn't give you star ratings.

- Workload Balancing does not automatically power off any hosts beyond the number specified in the Failures allowed box in the **Configure HA** dialog.
 - However, Workload Balancing might still make recommendations to power off more hosts than the number of host failures to tolerate. (For example, Workload Balancing still recommends that you power off two hosts when High Availability is only configured to tolerate one host failure.) However, when you attempt to apply the recommendation, XenCenter might display an error message stating that High Availability is no longer guaranteed.
 - When Workload Balancing runs in automated mode and has power management enabled, recommendations that exceed the number of tolerated host failures are ignored. In this situation, the Workload Balancing log shows a message that power-management recommendation wasn't applied because High Availability is enabled.

Generate workload reports

Workload Balancing captures performance data and can use this data to generate reports, known as Workload Reports, about your virtualized environment, including reports about hosts and VMs. The Workload Balancing reports can help you perform capacity planning, determine virtual server health, and evaluate how effective your configured threshold levels are.

You can use the Pool Health report to evaluate how effective your optimization thresholds are. While Workload

Balancing provides default threshold settings, you might need to adjust these defaults for them to provide value in

your environment. If you do not have the optimization thresholds adjusted to the correct level for your environment,

Workload Balancing recommendations might not be appropriate for your environment.

To run reports, you do not need to configure for Workload Balancing to make placement recommendations or move VMs.

However, you must configure the Workload Balancing component. Ideally, you must set critical

thresholds to values that

reflect the point at which the performance of the hosts in your pool degrades. Ideally, the pool has been running

Workload Balancing for a couple of hours or long enough to generate the data to display in the reports.

Workload Balancing lets you generate reports on three types of objects: physical hosts, resource pools, and VMs. At a

high level, Workload Balancing provides two types of reports:

- Historical reports that display information by date
- “Roll up” style reports, which provide a summarizing overview of an area
- Reports for auditing purposes, so you can determine, for example, the number of times a VM moved
- Chargeback report that shows VM usage and can help you measure and assign costs

Generate a Workload Balancing report

1. In XenCenter, from the **Pool** menu, select **View Workload Reports**.

You can also display the **Workload Reports** screen from the **WLB** tab by clicking the **Reports** button.

2. From the **Workload Reports** screen, select a report from the **Reports** pane.
3. Select the **Start Date** and the **End Date** for the reporting period. Depending on the report you select, you might be required to specify a host in the **Host** list.
4. Click **Run Report**. The report displays in the report window. For information about the meaning of the reports, see [Workload Balancing report glossary](#).

Navigate in a Workload Balancing report

After generating a report, you can use the toolbar buttons in the report to navigate and perform certain tasks. To

display the name of a toolbar button, pause your mouse over toolbar icon.

Toolbar buttons	Description
	Document Map enables you to display a document map that helps you navigate the report.
	Page Forward/Back enables you to move one page ahead or back in the report.
	Back to Parent Report enables you to return to the parent report when working in a sub-report.
	Stop Rendering cancels the report generation.
	Print enables you to print a report and specify general printing options. The report is printed in the default printer.
	Print Layout enables you to display a preview of the report before you print it.
	Page Setup enables you to specify printing options such as the paper size, page range, and page number.
	Export enables you to export the report as an Acrobat (.PDF) file or as an Excel (.XLS) file.
	Find enables you to search for a word in a report, such as the name of a VM.

Export a Workload Balancing report

You can export a report in either Microsoft Excel or Adobe Acrobat (PDF) formats.

1. After generating the report, click the following Export button:
2. Select one of the following items from the Export button menu:
 - Excel
 - Acrobat (PDF) file

Note:

Depending on the export format you select, the report contains different amounts of data. Reports exported to Excel include all the data available for reports, including “drilldown” data. Reports exported to PDF and displayed in XenCenter only contain the data that you selected when you generated the report.

Workload Balancing report glossary

This section provides information about the following Workload Balancing reports:

- [Chargeback Utilization Analysis](#)
- [Host Health History](#)

- [Pool Optimization Performance History](#)
- [Pool Audit Trail](#)
- [Pool Health](#)
- [Pool Health History](#)
- [Pool Optimization History](#)
- [Virtual Machine Motion History](#)
- [Virtual Machine Performance History](#)

Chargeback Utilization Analysis

You can use the Chargeback Utilization Analysis report (“chargeback report”) to determine how much of a resource a specific department in your organization used. Specifically, the report shows information about all the VMs in your pool, including their availability and resource utilization. Since this report shows VM up time, it can help you demonstrate Service Level Agreements compliance and availability.

The chargeback report can help you implement a simple chargeback solution and facilitate billing. To bill customers for a specific resource, generate the report, save it as Excel, and edit the spreadsheet to include your price per unit. Alternatively, you can import the Excel data into your billing system.

If you want to bill internal or external customers for VM usage, consider incorporating department or customer names in your VM naming conventions. This practice makes reading chargeback reports easier.

The resource reporting in the chargeback report is, sometimes, based on the allocation of physical resources to individual VMs.

The average memory data in this report is based on the amount of memory currently allocated to the VM. XenServer enables you to have a fixed memory allocation or an automatically adjusting memory allocation (Dynamic Memory Control).

The chargeback report contains the following columns of data:

- **VM Name.** The name of the VM to which the data in the columns in that row applies.
- **VM Uptime.** The number of minutes the VM was powered on (or, more specifically, appears with a green icon beside

it in XenCenter).

- **vCPU Allocation.** The number of virtual CPUs configured on the VM. Each virtual CPU receives an equal share of the physical CPUs on the host. For example, consider the case where you configured eight virtual CPUs on a host that contains two physical CPUs. If the **vCPU Allocation** column has “1” in it, this value is equal to 2/16 of the total processing power on the host.
- **Minimum CPU Usage (%).** The lowest recorded value for virtual CPU utilization in the reporting period. This value is expressed as a percentage of the VM’s vCPU capacity. The capacity is based on the number of vCPUs allocated to the VM. For example, if you allocated one vCPU to a VM, **Minimum CPU Usage** represents the lowest percentage of vCPU usage that is recorded. If you allocated two vCPUs to the VM, the value is the lowest usage of the combined capacity of both vCPUs as a percentage.

Ultimately, the percentage of CPU usage represents the lowest recorded workload that virtual CPU handled. For example, if you allocate one vCPU to a VM and the pCPU on the host is 2.4 GHz, 0.3 GHz is allocated to the VM. If the **Minimum CPU Usage** for the VM was 20%, the VM’s lowest usage of the physical host’s CPU during the reporting period was 60 MHz.
- **Maximum CPU Usage (%).** The highest percentage of the VM’s virtual CPU capacity that the VM consumed during the reporting period. The CPU capacity consumed is a percentage of the virtual CPU capacity you allocated to the VM.

For example, if you allocated one vCPU to the VM, the Maximum CPU Usage represents the highest recorded percentage of vCPU usage during the time reported. If you allocated two virtual CPUs to the VM, the value in this column represents the highest utilization from the combined capacity of both virtual CPUs.
- **Average CPU Usage (%).** The average amount, expressed as a percentage, of the VM’s virtual CPU capacity that was in use during the reporting period. The CPU capacity is the virtual CPU capacity you allocated to the VM. If you allocated two virtual CPUs to the VM, the value in this column represents the average utilization

from the combined
capacity of both virtual CPUs.

- **Total Storage Allocation (GB).** The amount of disk space that is currently allocated to the VM at the time the report was run. Frequently, unless you modified it, this disk space is the amount of disk space you allocated to the VM when you created it.
- **Virtual NIC Allocation.** The number of virtual interfaces (VIFs) allocated to the VM.
- **Current Minimum Dynamic Memory (MB).**
 - **Fixed memory allocation.** If you assigned a VM a fixed amount of memory (for example, 1,024 MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).
 - **Dynamic memory allocation.** If you configured XenServer to use Dynamic Memory Control, the minimum amount of memory specified in the range appears in this column. If the range has 1,024 MB as minimum memory and 2,048 MB as maximum memory, the **Current Minimum Dynamic Memory (MB)** column displays 1,024 MB.
- **Current Maximum Dynamic Memory (MB).**
 - **Dynamic memory allocation.** If XenServer adjusts a VM's memory automatically based on a range, the maximum amount of memory specified in the range appears in this column. For example, if the memory range values are 1,024 MB minimum and 2,048 MB maximum, 2,048 MB appears in the **Current Maximum Dynamic Memory (MB)** column.
 - **Fixed memory allocation.** If you assign a VM a fixed amount of memory (for example, 1,024 MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).
- **Current Assigned Memory (MB).**
 - **Dynamic memory allocation.** When Dynamic Memory Control is configured, this value indicates the amount of memory XenServer allocates to the VM when the report runs.

- **Fixed memory allocation.** If you assign a VM a fixed amount of memory (for example, 1,024 MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).

Note:

If you change the VM's memory allocation immediately before running this report, the value reflected in this column reflects the new memory allocation you configured.

- **Average Assigned Memory (MB).**

- **Dynamic memory allocation.** When Dynamic Memory Control is configured, this value indicates the average amount of memory XenServer allocated to the VM over the reporting period.
- **Fixed memory allocation.** If you assign a VM a fixed amount of memory (for example, 1,024 MB), the same amount of memory appears in the following columns: Current Minimum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).

Note:

If you change the VM's memory allocation immediately before running this report, the value in this column might not change from what was previously displayed. The value in this column reflects the average over the time period.

- **Average Network Reads (BPS).** The average amount of data (in bits per second) the VM received during the reporting period.
- **Average Network Writes (BPS).** The average amount of data (in bits per second) the VM sent during the reporting period.
- **Average Network Usage (BPS).** The combined total (in bits per second) of the Average Network Reads and Average Network Writes. If a VM sends, on average, 1,027 bps and receives, on average, 23,831 bps during the reporting period, the Average Network Usage is the combined total of these values: 24,858 bps.

- **Total Network Usage (BPS).** The total of all network read and write transactions in bits per second over the reporting period.

Host Health History

This report displays the performance of resources (CPU, memory, network reads, and network writes) on specific host in relation to threshold values.

The colored lines (red, green, yellow) represent your threshold values. You can use this report with the Pool Health report for a host to determine how the host's performance might affect overall pool health. When you are editing the performance thresholds, you can use this report for insight into host performance.

You can display resource utilization as a daily or hourly average. The hourly average lets you see the busiest hours of the day, averaged, for the time period.

To view report data which is grouped by hour, under **Host Health History** expand **Click to view report data grouped by hour for the time period.**

Workload Balancing displays the average for each hour for the time period you set. The data point is based on a utilization average for that hour for all days in the time period. For example, in a report for May 1, 2009, to May 15, 2009, the Average CPU Usage data point represents the resource utilization of all 15 days at 12:00 hours. This information is combined as an average. If CPU utilization was 82% at 12PM on May 1, 88% at 12PM on May 2, and 75% on all other days, the average displayed for 12PM is 76.3%.

Note:

Workload Balancing smoothes spikes and peaks so data does not appear artificially high.

Pool Optimization Performance History

The optimization performance report displays optimization events against that pool's average resource usage. These events are instances when you optimized a resource pool. Specifically, it displays resource usage for

CPU, memory,
network reads, and network writes.

The dotted line represents the average usage across the pool over the period of days you select. A blue bar indicates the day on which you optimized the pool.

This report can help you determine if Workload Balancing is working successfully in your environment. You can use this report to see what led up to optimization events (that is, the resource usage before Workload Balancing recommended optimizing).

This report displays average resource usage for the day. It does not display the peak utilization, such as when the system is stressed. You can also use this report to see how a resource pool is performing when Workload Balancing is not making optimization recommendations.

In general, resource usage declines or stays steady after an optimization event. If you do not see improved resource usage after optimization, consider readjusting threshold values. Also, consider whether the resource pool has too many VMs and whether you added or removed new VMs during the period that you specified.

Pool Audit Trail

This report displays the contents of the XenServer Audit Log. The Audit Log is a XenServer feature designed to log attempts to perform unauthorized actions and select authorized actions. These actions include:

- Import and export
- Host and pool backups
- Guest and host console access.

The report gives more meaningful information when you give XenServer administrators their own user accounts with distinct roles assigned to them by using the RBAC feature.

Important:

To run the audit log report, you must enable the Audit Logging feature. By default, Audit Log is

always enabled in
the Workload Balancing virtual appliance.

The enhanced Pool Audit Trail feature allows you to specify the granularity of the audit log report. You can also search and filter the audit trail logs by specific users, objects, and by time. The Pool Audit Trail Granularity is set to Minimum by default. This option captures limited amount of data for specific users and object types. You can modify the setting at any time based on the level of detail you require in your report. For example, set the granularity to Medium for a user-friendly report of the audit log. If you require a detailed report, set the option to Maximum.

The Pool Audit Trail report contains the following information:

- Time. The time XenServer recorded the user's action.
- User Name. The name of the person who created the session in which the action was performed. Sometimes, this value can be the User ID
- Event Object. The object that was the subject of the action (for example, a VM).
- Event Action. The action that occurred. For definitions of these actions, see [Audit Log Event Names](#).
- Access. Whether the user had permission to perform the action.
- Object Name. The name of the object (for example, the name of the VM).
- Object UUID. The UUID of the object (for example, the UUID of the VM).
- Succeeded. This information provides the status of the action (that is, whether it was successful).

Audit Log event names The Audit Log report logs XenServer events, event objects and actions, including import/export, host and pool backups, and guest and host console access. The following table defines some of the typical events that appear frequently in the XenServer Audit Log and Pool Audit Trail report. The table also specifies the granularity of these events.

In the Pool Audit Trail report, the events listed in the **Event Action** column apply to a pool, VM, or host. To

determine what the events apply to, see the [Event Object](#) and [Object Name](#) columns in the report. For more event definitions, see the events section of the [XenServer Management API](#).

Pool Audit Trail Granularity	Event Action	User Action
Minimum	<code>pool.join</code>	Instructed the host to join a new pool
Minimum	<code>pool.join_force</code>	Instructed (forced) the host to join a pool
Medium	<code>SR.destroy</code>	Destroyed the storage repository
Medium	<code>SR.create</code>	Created a storage repository
Medium	<code>VDI.snapshot</code>	Took a read-only snapshot of the VDI, returning a reference to the snapshot
Medium	<code>VDI.clone</code>	Took an exact copy of the VDI, returning a reference to the new disk
Medium	<code>VIF.plugin</code>	Hot-plugged the specified VIF, dynamically attaching it to the running VM
Medium	<code>VIF.unplug</code>	Hot-unplugged the specified VIF, dynamically detaching it from the running VM
Maximum	<code>auth.get_subject_identifier</code>	Queried the external directory service to obtain the subject identifier as a string from the human-readable subject name
Maximum	<code>task.cancel</code>	Requested that a task is canceled
Maximum	<code>VBD.insert</code>	Inserted new media into the device
Maximum	<code>VIF.get_by_uuid</code>	Obtained a reference to the VIF instance with the specified UUID
Maximum	<code>VDI.get_shareable</code>	Obtained the shareable field of the given VDI

Pool Audit Trail Granularity	Event Action	User Action
Maximum	<code>SR.get_all</code>	Returns a list of all SRs known to the system
Maximum	<code>pool.create_new_blob</code>	Created a placeholder for a named binary piece of data that is associated with this pool
Maximum	<code>host.send_debug_keys</code>	Injected the given string as debugging keys into Xen
Maximum	<code>VM.get_boot_record</code>	Returned a record describing the VMs dynamic state, initialized when the VM boots, and updated to reflect runtime configuration changes, for example, CPU hotplug

Pool Health

The Pool Health report displays the percentage of time a resource pool and its hosts spent in four different threshold ranges: Critical, High, Medium, and Low. You can use the Pool Health report to evaluate the effectiveness of your performance thresholds.

A few points about interpreting this report:

- Resource utilization in the Average Medium Threshold (blue) is the optimum resource utilization regardless of the placement strategy you selected. Likewise, the blue section on the pie chart indicates the amount of time that host used resources optimally.
- Resource utilization in the Average Low Threshold Percent (green) is not necessarily positive. Whether Low resource utilization is positive depends on your placement strategy. If your placement strategy is Maximum Density and resource usage is green, Workload Balancing might not be fitting the maximum number of VMs on that host or pool. If so, adjust your performance threshold values until most of your resource utilization falls into the Average Medium (blue) threshold range.

- Resource utilization in the Average Critical Threshold Percent (red) indicates the amount of time average resource utilization met or exceeded the Critical threshold value.

If you double-click on a pie chart for a host's resource usage, XenCenter displays the Host Health History report for that resource on that host. Clicking **Back to Parent Report** on the toolbar returns you to the Pool Health report.

If you find that most of your report results are not in the Average Medium Threshold range, adjust the Critical threshold for this pool. While Workload Balancing provides default threshold settings, these defaults are not effective in all environments. If you do not have the thresholds adjusted to the correct level for your environment, the Workload Balancing optimization and placement recommendations might not be appropriate. For more information, see [Change the critical thresholds](#).

Pool Health History

This report provides a line graph of resource utilization on all physical hosts in a pool over time. It lets you see the trend of resource utilization—if it tends to be increasing in relation to your thresholds (Critical, High, Medium, and Low). You can evaluate the effectiveness of your performance thresholds by monitoring trends of the data points in this report.

Workload Balancing extrapolates the threshold ranges from the values you set for the Critical thresholds when you connected the pool to Workload Balancing. Although similar to the Pool Health report, the Pool Health History report displays the average utilization for a resource on a specific date. Instead of the amount of time overall the resource spent in a threshold.

Except for the Average Free Memory graph, the data points never average above the Critical threshold line (red). For the Average Free Memory graph, the data points never average *below* the Critical threshold line (which is at the bottom of the graph). Because this graph displays *free* memory, the Critical threshold is a low value,

unlike the
other resources.

A few points about interpreting this report:

- When the Average Usage line in the chart approaches the Average Medium Threshold (blue) line, it indicates the pool's resource utilization is optimum. This indication is regardless of the placement strategy configured.
- Resource utilization approaching the Average Low Threshold (green) is not necessarily positive. Whether Low resource utilization is positive depends on your placement strategy. In the case where:
 - Your placement strategy is Maximum Density
 - Most days the Average Usage line is at or below the green lineWorkload Balancing might not be placing VMs as densely as possible on that pool. If so, adjust the pool's Critical threshold values until most of its resource utilization falls into the Average Medium (blue) threshold range.
- When the Average Usage line intersects with the Average Critical Threshold (red), it indicates when the average resource utilization met or exceeded the Critical threshold value for that resource.

If data points in your graphs aren't in the Average Medium Threshold range, but the performance is satisfactory, you can adjust the Critical threshold for this pool. For more information, see [Change the critical thresholds](#).

Pool Optimization History

The Pool Optimization History report provides chronological visibility into Workload Balancing optimization activity.

Optimization activity is summarized graphically and in a table. Drilling into a date field within the table displays detailed information for each pool optimization performed for that day.

This report lets you see the following information:

- VM Name: The name of the VM that Workload Balancing optimized.
- Reason: The reason for the optimization.
- Method: Whether the optimization was successful.

- From Host: The physical host where the VM was originally hosted.
- To Host: The physical host where the VM was migrated.
- Time: The time when the optimization occurred.

Tip:

You can also generate a Pool Optimization History report from the WLB tab by clicking the View History link.

Virtual Machine Motion History

This line graph displays the number of times VMs migrated on a resource pool over a period. It indicates if a migration resulted from an optimization recommendation and to which host the VM moved. This report also indicates the reason for the optimization. You can use this report to audit the number of migrations on a pool.

Some points about interpreting this report:

- The numbers on the left side of the chart correspond with the number of migrations possible. This value is based on how many VMs are in a resource pool.
- You can look at details of the migrations on a specific date by expanding the + sign in the Date section of the report.

Virtual Machine Performance History

This report displays performance data for each VM on a specific host for a time period you specify. Workload Balancing

bases the performance data on the amount of virtual resources allocated for the VM. For example, if Average CPU Usage for your VM is 67%, your VM uses, on average, 67% of its vCPU for the specified period.

The initial view of the report displays an average value for resource utilization over the period you specified.

Expanding the + sign displays line graphs for individual resources. You can use these graphs to see trends in resource utilization over time.

This report displays data for CPU Usage, Free Memory, Network Reads/Writes, and Disk Reads/Writes.

Configure Workload Balancing behavior

June 26, 2026

After connecting to the Workload Balancing virtual appliance, you can edit the settings Workload Balancing uses to calculate placement and recommendations. Workload Balancing settings apply collectively to all VMs and hosts in the pool.

Placement and optimization settings that you can modify include the following:

- Changing the placement strategy
- Configuring automatic optimizations and power management
- Editing performance thresholds and metric weightings
- Excluding hosts.

Provided the network and disk thresholds align with the hardware in your environment, consider using most of the defaults in Workload Balancing initially. After Workload Balancing is enabled for a while, we recommend evaluating your performance thresholds and determining whether to edit them. For example, consider the following cases:

- Getting recommendations when they are not yet required. If so, try adjusting the thresholds until Workload Balancing begins providing suitable recommendations.
- Not getting recommendations when you expect to receive them. For example, if your network has insufficient bandwidth and you do not receive recommendations, you might have to tweak your settings. If so, try lowering the network critical thresholds until Workload Balancing begins providing recommendations.

Before you edit your thresholds, you can generate a Pool Health report and the Pool Health History report for each physical host in the pool. For more information, see [Generate workload reports](#).

Notes:

- Workload Balancing is available for XenServer® Premium Edition customers. For more information about XenServer

licensing, see [Licensing](#). To upgrade, or to get a XenServer license, visit the [XenServer website](#).

- Workload Balancing 8.4.1 is compatible with XenServer 8.4 and XenServer 9.
- The base OS of the Workload Balancing virtual appliance was originally derived from a general-purpose Linux distribution. The Workload Balancing virtual appliance in its entirety is supported by XenServer as documented in the [XenServer Product Matrix](#).

This article assumes that you already connected your pool to a Workload Balancing virtual appliance. For information about downloading, importing, configuring, and connecting to a Workload Balancing virtual appliance, see [Get started](#).

Adjust the optimization mode

Workload Balancing makes recommendations to rebalance, or optimize, the VM workload in your environment based on a strategy for placement you select. The placement strategy is known as the optimization mode.

You can choose from the following optimization modes:

- **Maximize Performance** (default)

Workload Balancing attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, and network pressure for all hosts. When Maximize Performance is your placement strategy, Workload Balancing recommends optimization when a host reaches the High threshold.

- **Maximize Density**

Workload Balancing attempts to minimize the number of physical hosts that must be online by consolidating the active VMs.

When you select Maximize Density as your placement strategy, you can specify parameters similar to the ones in Maximize Performance. However, Workload Balancing uses these parameters to determine how it can pack VMs onto a host. If Maximize Density is your placement strategy, Workload Balancing recommends consolidation optimizations when a VM reaches the Low threshold.

Workload Balancing also lets you apply these optimization modes always, *fixed*, or switch between modes for specified time periods, *scheduled*:

Fixed optimization modes

Fixed optimization modes set Workload Balancing to have a specific optimization behavior always. This behavior can be either to try to create the best performance or to create the highest density.

To set a fixed optimization mode, complete the following steps:

1. In XenCenter®, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, click **Optimization Mode**.
5. In the **Fixed** section of the **Optimization Mode** page, select one of these optimization modes:
 - **Maximize Performance** (default). Attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, and network pressure for all hosts.
 - **Maximize Density**. Attempts to fit as many VMs as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online.

Scheduled optimization modes

Scheduled optimization modes let you schedule for Workload Balancing to apply different optimization modes depending on the time of day. For example, you might want to configure Workload Balancing to optimize for performance during the day when you have users connected. To save energy, you can then specify for Workload Balancing to optimize for Maximum Density at night.

When you configure scheduled optimization modes, Workload Balancing automatically changes to the optimization mode at the beginning of the time period you specified. You can configure Everyday, Weekdays, Weekends, or individual days.

For the hour, you select a time of day.

To set a schedule for your optimization modes, complete the following steps:

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, click **Optimization Mode**.
5. In the **Optimization Mode** pane, select **Scheduled**. The **Scheduled** section becomes available.
6. Click **Add New**.
7. In the **Change to** box, select one of the following modes:
 - **Maximize Performance**. Attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, and network pressure for all hosts.
 - **Maximize Density**. Attempts to fit as many VMs as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online.
8. Select the day of the week and the time when you want Workload Balancing to begin operating in this mode.
9. Repeat the preceding steps to create more scheduled mode tasks until you have the number you need. If you only schedule one task, Workload Balancing switches to that mode as scheduled, but then it never switches back.
10. Click **OK**.

To change your schedule settings, complete the following steps.

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, click **Optimization Mode**.
5. Select the task that you want to delete or disable from the **Scheduled Mode Changes** list.
6. Do one of the following:
 - **Delete the task permanently**: Click the **Delete** button.
 - **Stop the task from running temporarily**: Right-click the task and click **Disable**.

Tips:

- You can also disable or enable tasks by selecting the task, clicking **Edit**, and selecting the **Enable Task** check box in the **Optimization Mode Scheduler** dialog.
- To re-enable a task, right-click the task in the **Scheduled Mode Changes** list and click **Enable**.

- **Edit the task:** Double-click the task that you want to edit. In the **Change to** box, select a different mode or make other changes as desired.

Note:

Clicking Cancel, before clicking OK, undoes any changes you made in the Optimization tab, including deleting a task.

Optimize and manage power automatically

You can configure Workload Balancing to apply recommendations automatically and turn hosts on or off automatically. To power down hosts automatically (for example, during low-usage periods), you must configure Workload Balancing to apply recommendations automatically and enable power management. Both power management and automation are described in the sections that follow.

Apply recommendations automatically

Workload Balancing lets you configure for it to apply recommendations on your behalf and perform the optimization actions it recommends automatically. You can use this feature, which is known as automatic optimization acceptance, to apply any recommendations automatically, including ones to improve performance or power down hosts. However, to power down hosts as VMs usage drops, you must configure automation, power management, and Maximum Density mode.

By default, Workload Balancing does not apply recommendations automatically. If you want Workload Balancing to apply

recommendations automatically, enable automation. If you do not, you must apply recommendations manually by clicking

Apply Recommendations.

Workload Balancing does not automatically apply recommendations to hosts or VMs when the recommendations conflict with

HA settings. If a pool becomes overcommitted by applying Workload Balancing optimization recommendations, XenCenter

prompts you whether you want to continue applying the recommendation. When automation is enabled, Workload Balancing

does not apply any power-management recommendations that exceed the number of host failures to tolerate in the HA

plan.

When Workload Balancing is running with the automation feature enabled, this behavior is sometimes called running in automated mode.

It is possible to tune how Workload Balancing applies recommendations in automated mode. For information, see

[Set conservative or aggressive automated recommendations.](#)

To apply optimization recommendations automatically

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, click **Automation**.
5. Select one or more of the following check boxes:
 - **Automatically apply Optimization recommendations.** When you select this option, you do not need to accept optimization recommendations manually. Workload Balancing automatically accepts the optimization and placement recommendations that it makes.
 - **Automatically apply Power Management recommendations.** The behavior of this option varies according to the optimization mode of the pool:
 - Maximum Performance mode: When **Automatically apply Power Management recommendations** is enabled,

Workload Balancing automatically powers on hosts when doing so improves host performance.

- **Maximum Density mode:** When **Automatically apply Power Management recommendations** is enabled, Workload Balancing automatically powers off hosts when resource utilization drops below the Low threshold. That is, Workload Balancing powers off hosts automatically during low usage periods.

6. (Optional) Fine-tune optimization recommendations by clicking **Advanced** in the left pane of the **Settings** dialog and doing one or more of the following actions:

- Specifying the number of times Workload Balancing must make an optimization recommendation before the recommendation is applied automatically. The default is three times, which means the recommendation is applied on the third time it is made.
- Selecting the lowest level of optimization recommendation that you want Workload Balancing to apply automatically. The default is High.
- Changing the aggressiveness with which Workload Balancing applies its optimization recommendations.

You might also want to specify the number of minutes Workload Balancing has to wait before applying an optimization recommendation to a recently moved VM.

All of these settings are explained in more depth in [Set conservative or aggressive automated recommendations](#).

7. (optional) If you want to configure power management, click **Automation/Power Management**

- a) In the **Power Management** section, select the hosts that you want Workload Balancing to recommend powering on and off.

Note:

Selecting hosts for power management recommendations without selecting **Automatically apply Power Management recommendations** causes Workload Balancing to suggest power-

management recommendations but not apply them automatically for you.

If none of the hosts in the resource pool support remote power management, Workload Balancing displays the message, “No hosts support Power Management.”

b) Click **OK**.

8. To finish configuring the automation, click **OK**.

Enable Workload Balancing power management

The term power management means the ability to turn the power on or off for physical hosts. In a Workload Balancing context, this term means powering hosts in a pool on or off based on the total workload of the pool.

Configuring Workload Balancing power management on a host requires that:

- The hardware for the host has remote power on/off capabilities.
- The Host Power On feature is configured for the host. To configure the Host Power On feature for the host, see [Configure Host Power On feature](#).
- The host has been explicitly selected as a host to participate in Workload Balancing power management.

In addition, if you want Workload Balancing to power off hosts automatically, configure Workload Balancing to do the following actions:

- Apply recommendations automatically
- Apply power management recommendations automatically

When a host is set to participate in power management, Workload Balancing makes power-on and power-off recommendations as needed.

If you run in Maximum Density mode:

- When Workload Balancing detects unused resources in a pool, it recommends powering off hosts until it eliminates all excess capacity.

- If there isn't enough host capacity in the pool to shut down hosts, Workload Balancing recommends leaving the hosts on until the pool workload decreases enough.
- When you configure Workload Balancing to power off extra hosts automatically, it applies these recommendations automatically and, so, behaves in the same way.

If you run in Maximum Performance mode:

- If you configure Workload Balancing to power on hosts automatically, Workload Balancing powers on hosts when resource utilization on a host exceeds the High threshold.
- Workload Balancing never powers off hosts after it has powered them on.

If you turn on the option to apply power management recommendations automatically, you do so at the pool level.

However, you can specify which hosts from the pool you want to participate in power management.

Configure Host Power On feature To configure the Host Power On feature for your host, follow these steps:

1. In XenCenter, select your host and click **Properties**.
2. In the left pane, click **Power On**.
3. For the **Power On mode**, select a Power On mode which can enable remote server startup.
4. After a Power On mode is configured, select your pool.
5. In the **Properties** pane of the pool, click the **WLB** tab.
6. In the **WLB** tab, click **Settings**.
7. In the left pane, click **Automation**.
8. For **Automation**, select the following check boxes:
 - **Automatically apply Optimization recommendations.** When you select this option, you do not need to accept optimization recommendations manually. Workload Balancing automatically accepts the optimization and placement recommendations that it makes.
 - **Automatically apply Power Management recommendations.** The behavior of this option varies according to the optimization mode of the pool:

- Maximum Performance mode: When **Automatically apply Power Management recommendations** is enabled, Workload Balancing automatically powers on hosts when doing so improves host performance.
- Maximum Density mode: When **Automatically apply Power Management recommendations** is enabled, Workload Balancing automatically powers off hosts when resource utilization drops below the Low threshold. That is, Workload Balancing powers off hosts automatically during low usage periods.

9. For **Power Management**, select the name of the **Host Server** that you're currently configuring.


Understand power management behavior

Before Workload Balancing recommends powering hosts on or off, it selects the hosts to transfer VMs to. It does so in the following order:

1. Filling the pool coordinator since it is the host that cannot be powered off.
2. Filling the host with the most VMs.
3. Filling subsequent hosts according to which hosts have the most VMs running.

When Workload Balancing fills the pool coordinator, it does so assuming artificially low thresholds for the coordinator. Workload Balancing uses these low thresholds as a buffer to prevent the pool coordinator from being overloaded.

Workload Balancing fills hosts in this order to encourage density.



When Workload Balancing detects a performance issue while the pool is in Maximum Density mode, it recommends migrating workloads among the powered-on hosts. If Workload Balancing cannot resolve the issue using this method, it attempts to power on a host. Workload Balancing determines which hosts to power on by applying the same criteria that it would if the optimization mode was set to Maximum Performance.

When Workload Balancing runs in Maximum Performance mode, Workload Balancing recommends powering on hosts until the resource utilization on all pool members falls below the High threshold.

While migrating VMs, if Workload Balancing determines that increasing capacity benefits the overall performance of the

pool, it powers on hosts automatically or recommends doing so.

Important:

Workload Balancing only recommends powering on a host that Workload Balancing powered off.

Design environments for power management and VM consolidation

When you are planning a XenServer implementation and you intend to configure automatic VM consolidation and power management, consider your workload design. For example, you might want to:

- Place different types of workloads in separate pools.

If you have an environment with distinct types of workloads, consider whether to locate the VMs hosting these workloads in different pools. Also consider splitting VMs that host types of applications that perform better with certain types of hardware into different pools.

Because power management and VM consolidation are managed at the pool level, design pools so they contain workloads that you want consolidated at the same rate. Ensure that you factor in considerations such as those discussed in [Configure advanced settings](#).

- Exclude hosts from Workload Balancing.

Some hosts might need to be always on. For more information, see [Exclude hosts from recommendations](#).

- Exclude VMs from Workload Balancing.

Some VMs might need to be excluded from Workload Balancing optimization and placement recommendations. For more information, see [Exclude VMs from recommendations](#).

Understand when Workload Balancing makes recommendations

Workload Balancing continuously evaluates the resource metrics of physical hosts and VMs across the pools that it is managing against thresholds. Thresholds are preset values that function like boundaries that a host must exceed before

Workload Balancing can make an optimization recommendation. The Workload Balancing process is as follows:

1. Workload Balancing detects that the threshold for a resource was violated.
2. Workload Balancing evaluates if it makes an optimization recommendation.
3. Workload Balancing determines which hosts it recommends function as the destination hosts and in what order to make any optimizations. A destination host is the host where Workload Balancing recommends relocating one or more VMs.
4. Workload Balancing makes an optimization recommendation.

When evaluating hosts in the pool to make an optimization recommendation, Workload Balancing uses thresholds and weightings as follows:

- **Thresholds** are the boundary values that Workload Balancing compares the resource metrics of your pool against.
The thresholds are used to determine whether to make a recommendation and what hosts are a suitable candidate for hosting relocated VMs.
- **Weightings** are a way of ranking resources according to how much you want them to be considered, are used to determine the processing order. After Workload Balancing decides to make a recommendation, it uses your specifications of which resources are important to determine the following:
 - Which hosts' performance to address first
 - Which VMs to recommend migrating first

For each resource Workload Balancing monitors, it has four levels of thresholds: Critical, High, Medium, and Low.

Workload Balancing evaluates whether to make a recommendation when a resource metric on a host:

- Exceeds the High threshold when the pool is running in Maximum Performance mode (improve performance)
- Drops below the Low threshold when the pool is running in Maximum Density mode (consolidate VMs on hosts)
- Exceeds the Critical threshold when the pool is running in Maximum Density mode (improve performance)

If the High threshold for a pool running in Maximum Performance mode is 80%, when CPU utilization on a host reaches

80.1%, Workload Balancing evaluates whether to issue a recommendation.

When a resource violates its threshold, Workload Balancing evaluates the resource metric against historical

performance to prevent making an optimization recommendation based on a temporary spike. To do so, Workload Balancing

creates a historically averaged utilization metric by evaluating the data for resource utilization captured at the

following times:

Data captured	Weight
Immediately, at the time threshold was exceeded. That is, real-time data.	70%
30 minutes before the threshold was exceeded	25%
24 hours before the threshold was exceeded	5%

If CPU utilization on the host exceeds the threshold at 12:02 PM, Workload Balancing checks the utilization at 11:32

AM that day, and at 12:02PM on the previous day. For example, if CPU utilization is at the following values, Workload

Balancing doesn't make a recommendation:

- 80.1% at 12:02 PM that day
- 50% at 11:32 AM that day
- 78% at 12:02 PM the previous day

This behavior is because the historically averaged utilization is 72.5%, so Workload Balancing assumes that the

utilization is a temporary spike. However, if the CPU utilization was 83% at 11:32AM, Workload Balancing makes a

recommendation since the historically averaged utilization is 80.1%.

Optimization and consolidation process

The Workload Balancing process for determining potential optimizations varies according to the optimization mode -

Maximum Performance or Maximum Density. However, regardless of the optimization mode, the optimization and placement recommendations are made using a two-stage process:

1. Determine potential optimizations: which VMs to migrate off hosts.
2. Determine placement recommendations: which hosts would be suitable candidates for new VMs.

Note:

Workload Balancing only recommends migrating VMs that meet the XenServer criteria for live migration. One of these criteria is that the destination host must have the storage the VM requires. The destination host must also have sufficient resources to accommodate adding the VM without exceeding the thresholds of the optimization mode configured on the pool. For example, the High threshold in Maximum Performance mode and the Critical threshold for Maximum Density mode.

When Workload Balancing is running in automated mode, you can tune the way it applies recommendations. For more information, see [Set conservative or aggressive automated recommendations](#).

Optimization recommendation process in Maximum Performance mode When running in Maximum Performance mode, Workload Balancing uses the following process to determine potential optimizations:

1. Every two minutes Workload Balancing evaluates the resource utilization for each host in the pool. It does so by monitoring on each host and determining if each resource's utilization exceeds its High threshold. For more information, see [Change the critical threshold](#).

In Maximum Performance mode, if a utilization of a resource exceeds its High threshold, Workload Balancing starts the process to determine whether to make an optimization recommendation. Workload Balancing determines whether to make an optimization recommendation based on whether doing so can ease performance constraints, such as ones revealed by the High threshold.

For example, consider the case where Workload Balancing sees that insufficient CPU resources negatively affect the performance of the VMs on a host. If Workload Balancing can find another host with less CPU utilization, it recommends moving one or more VMs to another host.

2. If a resource's utilization on a host exceeds the relevant threshold, Workload Balancing combines the following data to form the historically averaged utilization:

- The resource's current utilization
- Historical data from 30 minutes ago
- Historical data from 24 hours ago

If the historically averaged utilization exceeds the threshold of a resource, Workload Balancing determines it makes an optimization recommendation.

3. Workload Balancing uses metric weightings to determine what hosts to optimize first. The resource to which you have assigned the most weight is the one that Workload Balancing attempts to address first. For more information, see [Tune metric weightings](#).

4. Workload Balancing determines which hosts can support the VMs it wants to migrate off hosts. Workload Balancing makes this determination by calculating the projected effect on resource utilization of placing different combinations of VMs on hosts. Workload Balancing uses a method of performing these calculations that in mathematics is known as permutation.

To do so, Workload Balancing creates a single metric or score to forecast the impact of migrating a VM to the host. The score indicates the suitability of a host as a home for more VMs.

To score host performance, Workload Balancing combines the following metrics:

- The current metrics of the host
- The metrics of the host from the last 30 minutes
- The metrics of the host from 24 hours ago
- The metrics of the VM.

5. After scoring hosts and VMs, Workload Balancing attempts to build virtual models of what the hosts look like with

different combinations of VMs. Workload Balancing uses these models to determine the best host to place the VM.

In Maximum Performance mode, Workload Balancing uses metric weightings to determine what hosts to optimize first and what VMs on those hosts to migrate first. Workload Balancing bases its models on the metric weightings. For example, if CPU utilization is assigned the highest importance, Workload Balancing sorts hosts and VMs to optimize according to the following criteria:

- a) What hosts are running closest to the High threshold for CPU utilization.
 - b) What VMs have the highest CPU utilization or are running the closest to its High threshold.
6. Workload Balancing continues calculating optimizations. It views hosts as candidates for optimization and VMs as candidates for migration until predicted resource utilization on the host hosting the VM drops below the High threshold. Predicted resource utilization is the resource utilization that Workload Balancing forecasts a host has after Workload Balancing has added or removed a VM from the host.

Consolidation process in Maximum Density mode Workload Balancing determines whether to make a recommendation based on whether it can migrate a VM onto a host and still run that host below the Critical threshold.

1. When a resource's utilization drops below its Low threshold, Workload Balancing begins calculating potential consolidation scenarios.
2. When Workload Balancing discovers a way that it can consolidate VMs on a host, it evaluates whether the destination host is a suitable home for the VM.
3. Like in Maximum Performance mode, Workload Balancing scores the host to determine the suitability of a host as a home for new VMs.

Before Workload Balancing recommends consolidating VMs on fewer hosts, it checks that resource utilization on those hosts after VMs are relocated to them is below Critical thresholds.

Note:

Workload Balancing does not consider metric weightings when it makes a consolidation recommendation. It only considers metric weightings to ensure performance on hosts.

4. After scoring hosts and VMs, Workload Balancing attempts to build virtual models of what the hosts look like with different combinations of VMs. It uses these models to determine the best host to place the VM.
5. Workload Balancing calculates the effect of adding VMs to a host until it forecasts that adding another VM causes a host resource to exceed the Critical threshold.
6. Workload Balancing recommendations always suggest filling the pool coordinator first since it is the host that cannot be powered off. However, Workload Balancing applies a buffer to the pool coordinator so that it cannot be over-allocated.
7. Workload Balancing continues to recommend migrating VMs on to hosts until all remaining hosts exceed a Critical threshold when a VM is migrated to them.

Change the critical thresholds

You might want to change critical thresholds as a way of controlling when optimization recommendations are triggered.

This section provides guidance about:

- How to modify the default Critical thresholds on hosts in the pool
- How values set for Critical threshold alter High, Medium, and Low thresholds.

Workload Balancing determines whether to produce recommendations based on whether the averaged historical utilization for a resource on a host violates its threshold. Workload Balancing recommendations are triggered when the High threshold in Maximum Performance mode or Low and Critical thresholds for Maximum Density mode are violated. For more information, see [Optimization and consolidation process](#).

After you specify a new Critical threshold for a resource, Workload Balancing resets the other thresholds of the

resource relative to the new Critical threshold. To simplify the user interface, the Critical threshold is the only threshold you can change through XenCenter.

The following table shows the default values for the Workload Balancing thresholds:

Metric	Critical	High	Medium	Low
CPU Utilization	90%	76.5%	45%	22.5%
Free Memory	51 MB	63.75 MB	510 MB	1020 MB
Network Reads	25 MB/sec	21.25 MB/sec	12.5 MB/sec	6.25 MB/sec
Network Writes	25 MB/sec	21.25 MB/sec	12.5 MB/sec	6.25 MB/sec
Disk Reads	25 MB/sec	21.25 MB/sec	12.5 MB/sec	6.25 MB/sec
Disk Writes	25 MB/sec	21.25 MB/sec	12.5 MB/sec	6.25 MB/sec

To calculate the threshold values for all metrics except memory, Workload Balancing multiplies the new value for the Critical threshold with the following factors:

- **High Threshold Factor:** 0.85
- **Medium Threshold Factor:** 0.50
- **Low Threshold Factor:** 0.25

For example, if you increase the Critical threshold for CPU utilization to 95%, Workload Balancing resets the other thresholds as follows:

- High: 80.75%
- Medium: 47.5%
- Low: 23.75%

To calculate the threshold values for free memory, Workload Balancing multiplies the new value for the Critical threshold with these factors:

- **High Threshold Factor:** 1.25
- **Medium Threshold Factor:** 10.0
- **Low Threshold Factor:** 20.0

For example, if you increase the Critical threshold for free memory to 45 MB, Workload Balancing resets the other thresholds as follows:

- High: 56.25 MB
- Medium: 450 MB
- Low: 900 MB

To perform this calculation for a specific threshold, multiply the factor for the threshold with the value you entered

for the critical threshold for that resource:

$$1 \text{ High, Medium, or Low Threshold} = \text{Critical Threshold} * \text{High, Medium, or Low Threshold Factor}$$

While the Critical threshold triggers many optimization recommendations, other thresholds can also trigger

optimization recommendations, as follows:

- **High threshold.**

- **Maximum Performance.** Exceeding the High threshold triggers optimization recommendations to relocate a VM to a host with lower resource utilization.
- **Maximum Density.** Workload Balancing doesn't recommend placing a VM on host when moving that VM to the host causes the host resource utilization to exceed a High threshold.

- **Low threshold.**

- **Maximum Performance.** Workload Balancing does not trigger recommendations from the Low threshold.
- **Maximum Density.** When a metric value drops below the Low threshold, Workload Balancing determines that hosts are underutilized and makes an optimization recommendation to consolidate VMs on fewer hosts. Workload Balancing continues to recommend moving VMs onto a host until the metric values for one of the host's resources reaches its High threshold.

However, after a VM is relocated, utilization of a resource on the VM's new host can exceed a Critical

threshold. In this case, Workload Balancing temporarily uses an algorithm similar to the Maximum Performance

load-balancing algorithm to find a new host for the VMs. Workload Balancing continues to use this algorithm to

recommend moving VMs until resource utilization on hosts across the pool falls below the High threshold.

To change the critical thresholds:

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, select **Critical Thresholds**. These critical thresholds are used to evaluate host resource utilization.
5. In the **Critical Thresholds** page, type one or more new values in the **Critical Thresholds** boxes. The values represent resource utilization on the host.

Workload Balancing uses these thresholds when making VM placement and pool-optimization recommendations. Workload Balancing strives to keep resource utilization on a host below the critical values set.

Tune metric weightings

How Workload Balancing uses metric weightings when determining which hosts and VMs to process first varies according to the optimization mode: Maximum Density or Maximum Performance. In general, metric weightings are used when a pool is in Maximum Performance mode. However, when Workload Balancing is in Maximum Density mode, it does use metric weightings when a resource exceeds its Critical threshold.

When Workload Balancing is processing optimization recommendations, it creates an optimization order. Workload Balancing determines the order by ranking the hosts according to which hosts have the highest metric values for whatever resource is ranked as the most important in the metric weightings page.

Maximum Performance mode

In Maximum Performance mode, Workload Balancing uses metric weightings to determine:

- On which hosts to first address the performance
- Which VMs to recommend migrating first

For example, if Network Writes is the most important resource, Workload Balancing first makes optimization recommendations for the host with the highest number of Network Writes per second. To make Network Writes the most important resource move the **Metric Weighting** slider to the right and all the other sliders to the middle.

If you configure all resources to be equally important, Workload Balancing addresses CPU utilization first and memory second, as these resources are typically the most constrained. To make all resources equally important, set the **Metric Weighting** slider in the same place for all resources.

Maximum Density mode

In Maximum Density mode, Workload Balancing only uses metric weightings when a host reaches the Critical threshold. At that point, Workload Balancing applies an algorithm similar to the algorithm for Maximum Performance until no hosts exceed the Critical thresholds. When using this algorithm, Workload Balancing uses metric weightings to determine the optimization order in the same way as it does for Maximum Performance mode.

If two or more hosts have resources exceeding their Critical thresholds, Workload Balancing verifies the importance you set for each resource. It uses this importance to determine which host to optimize first and which VMs on that host to relocate first.

For example, your pool contains host A and host B, which are in the following state:

- The CPU utilization on host A exceeds its Critical threshold and the metric weighting for CPU utilization is set to **More Important**.
- The memory utilization on host B exceeds its Critical threshold and the metric weighting for memory utilization is set to **Less Important**.

Workload Balancing recommends optimizing host A first because the resource on it that reached the Critical threshold is the resource assigned the highest weight. After Workload Balancing determines that it must address the performance

on host A, Workload Balancing then begins recommending placements for VMs on that host. It begins with the VM that has

the highest CPU utilization, since that CPU utilization is the resource with the highest weight.

After Workload Balancing has recommended optimizing host A, it makes optimization recommendations for host B. When it

recommends placements for the VMs on host B, it does so by addressing CPU utilization first, since CPU utilization was

assigned the highest weight. If there are more hosts that need optimization, Workload Balancing addresses the

performance on those hosts according to what host has the third highest CPU utilization.

By default, all metric weightings are set to the farthest point on the slider: More Important.

Note:

The weighting of metrics is relative. If all metrics are set to the same level, even if that level is Less

Important, they are all weighted the same. The relation of the metrics to each other is more important than the actual weight at which you set each metric.

To edit metric weighting factors

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, select **Metric Weighting**.
5. In **Metric Weighting** page, as desired, adjust the sliders beside the individual resources.

Move the slider towards **Less Important** to indicate that ensuring VMs always have the highest available amount of this resource is not as vital for this pool.

Exclude hosts from recommendations

When configuring Workload Balancing, you can specify that specific physical hosts are excluded from Workload Balancing

optimization and placement recommendations, including Start On placement recommendations.

Situations when you might want to exclude hosts from recommendations include when:

- You want to run the pool in Maximum Density mode and consolidate and shut down hosts, but you want to exclude specific hosts from this behavior.
- You have two VM workloads that must always run on the same host. For example, if the VMs have complementary applications or workloads.
- You have workloads that you do not want moved: for example, a domain controller or database server.
- You want to perform maintenance on a host and you do not want VMs placed on the host.
- The performance of the workload is so critical that the cost of dedicated hardware is irrelevant.
- Specific hosts are running high-priority workloads, and you do not want to use the HA feature to prioritize these VMs.
- The hardware in the host is not the optimum for the other workloads in the pool.

Regardless of whether you specify a fixed or scheduled optimization mode, excluded hosts remain excluded even when the optimization mode changes. Therefore, if you only want to prevent Workload Balancing from shutting off a host automatically, consider disabling Power Management for that host instead. For more information, see [Optimize and manage power automatically](#).

When you exclude a host from recommendations, you are specifying for Workload Balancing not to manage that host at all. This configuration means that Workload Balancing doesn't make any optimization recommendations for an excluded host. In contrast, when you don't select a host to participate in Power Management, Workload Balancing manages the host, but doesn't make power management recommendations for it.

To exclude hosts from Workload Balancing

Use this procedure to exclude a host in a pool that Workload Balancing is managing from power management, host evacuation, placement, and optimization recommendations.

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, select **Excluded Hosts**.

5. In **Excluded Hosts** page, select the hosts for which you do not want Workload Balancing to recommend alternate placements and optimizations.

Exclude VMs from recommendations

When configuring Workload Balancing, you can specify that specific VMs are excluded from Workload Balancing optimization and placement recommendations.

Situations when you might want to exclude VMs from recommendations include when:

- When two VM workloads must always run on the same host (for example, if they have complementary applications or workloads).
- You have workloads that you do not want moved (for example, domain controllers, DNS server, or SQL server).

Regardless of whether you specify a fixed or scheduled optimization mode, excluded VMs remain excluded even when the optimization mode changes.

To exclude VMs from placement and optimization recommendations

1. Select the VM in the **Resources** pane, select the **General** tab, and select **Properties**.
2. In the left pane, select **General**.
3. In the **General** page, in the **Tags:** field, select **Edit tags**.
4. Specify the new tag `wlb_excluded` and click **Create**.

Configure advanced settings

Workload Balancing supplies some advanced settings that let you control how Workload Balancing applies automated recommendations. These settings appear on the **Advanced** page of the Workload Balancing Configuration dialog. To get to the **Advanced** page, complete the following steps:

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.

4. In the left pane, select **Advanced**.

The following sections describe the behaviors that can be configured in the **Advanced** settings.

Set conservative or aggressive automated recommendations

When running in automated mode, the frequency of optimization and consolidation recommendations and how soon they are automatically applied is a product of multiple factors, including:

- How long you specify Workload Balancing waits after moving a VM before making another recommendation
- The number of recommendations Workload Balancing must make before applying a recommendation automatically
- The severity level a recommendation must achieve before the optimization is applied automatically
- The level of consistency in recommendations (recommended VMs to move, destination hosts) Workload Balancing requires before applying recommendations automatically

In general, only adjust the settings for these factors in the following cases:

- You have guidance from XenServer Technical Support
- You made significant observation and testing of the behavior of your pool with Workload Balancing enabled

Incorrectly configuring these settings can result in Workload Balancing not making any recommendations.

VM migration interval

You can specify the number of minutes Workload Balancing waits after the last time a VM was moved, before Workload

Balancing can make another recommendation for that VM. The recommendation interval is designed to prevent Workload

Balancing from generating recommendations for artificial reasons, for example, if there was a temporary utilization spike.

When automation is configured, it is especially important to be careful when modifying the recommendation interval. If

an issue occurs that leads to continuous, recurring spikes, decreasing the interval can generate many

recommendations
and, therefore, relocations.

Note:

Setting a recommendation interval does not affect how long Workload Balancing waits to factor recently rebalanced hosts into recommendations for Start-On Placement, Resume, and Maintenance Mode.

Recommendation count

Every two minutes, Workload Balancing checks to see if it can generate recommendations for the pool it is monitoring.

When you enable automation, you can specify the number of times a consistent recommendation must be made before

Workload Balancing automatically applies the recommendation. To do so, you configure a setting known as the

Recommendation Count, as specified in the **Recommendations** field. The **Recommendation Count** and the

Optimization Aggressiveness setting let you fine-tune the automated application of recommendations in your environment.

Workload Balancing uses the similarity of recommendations to make the following checks:

1. Whether the recommendation is truly needed
2. Whether the destination host has stable enough performance over a prolonged period to accept a relocated VM without needing to move it off the host again shortly

Workload Balancing uses the Recommendation Count value to determine whether a recommendation must be repeated before

Workload Balancing automatically applies the recommendation. Workload Balancing uses this setting as follows:

1. Every time Workload Balancing generates a recommendation that meets its consistency requirements, as indicated by the Optimization Aggressiveness setting, Workload Balancing increments the Recommendation Count. If the recommendation does not meet the consistency requirements, Workload Balancing might reset the Recommendation Count to zero. This behavior depends on the factors described in [Optimization aggressiveness](#).

2. When Workload Balancing generates enough consistent recommendations to meet the value for the Recommendation Count, as specified in the **Recommendations** field, it automatically applies the recommendation.

If you choose to modify this setting, the value to set varies according to your environment. Consider these scenarios:

- If host loads and activity increase rapidly in your environment, you might want to increase value for the Recommendation Count. Workload Balancing generates recommendations every two minutes. For example, if you set this interval to **3**, then six minutes later Workload Balancing applies the recommendation automatically.
- If host loads and activity increase gradually in your environment, you might want to decrease the value for the Recommendation Count.

Accepting recommendations uses system resources and affects performance when Workload Balancing is relocating the VMs.

Increasing the Recommendation Count increases the number of matching recommendations that must occur before Workload

Balancing applies the recommendation. This setting encourages Workload Balancing to apply more conservative, stable

recommendations and can decrease the potential for spurious VM moves. The Recommendation Count is set to a

conservative value by default.

Because of the potential impact adjusting this setting can have on your environment, only change it with extreme

caution. Preferably, make these adjustments by testing and iteratively changing the value or under the guidance of

XenServer Technical Support.

Recommendation severity

All optimization recommendations include a severity rating (Critical, High, Medium, Low) that indicates the importance

of the recommendation. Workload Balancing bases this rating on a combination of factors including the following:

- Configuration options you set, such as thresholds and metric tunings

- Resources available for the workload
- Resource-usage history.

The severity rating for a recommendation appears in the **Optimization Recommendations** pane on the **WLB** tab.

When you configure Workload Balancing to apply recommendations automatically, you can set the minimum severity level to associate with a recommendation before Workload Balancing automatically applies it.

Optimization aggressiveness

To provide extra assurance when running in automated mode, Workload Balancing has consistency criteria for accepting optimizations automatically. This criteria can help to prevent moving VMs due to spikes and anomalies. In automated mode, Workload Balancing does not accept the first recommendation it produces. Instead, Workload Balancing waits to apply a recommendation automatically until a host or VM exhibits consistent behavior over time. Consistent behavior over time includes factors like whether a host continues to trigger recommendations and whether the same VMs on that host continue to trigger recommendations.

Workload Balancing determines if behavior is consistent by using criteria for consistency and by having criteria for the number of times the same recommendation is made. You can configure how strictly you want Workload Balancing to apply the consistency criteria using the **Optimization Aggressiveness** setting. You can use this setting to control the amount of stability you want in your environment before Workload Balancing applies an optimization recommendation.

The most stable setting, Low aggressiveness, is configured by default. In this context, the term stable means the similarity of the recommended changes over time, as explained throughout this section. Aggressiveness is not desirable in most environments. Therefore, Low is the default setting.

Workload Balancing uses up to four criteria to ascertain consistency. The number of criteria that must be met varies according to the level you set in the **Optimization Aggressiveness** setting. The lower the level (for example, Low

or Medium) the less aggressive Workload Balancing is in accepting a recommendation. In other words, Workload Balancing

is stricter about requiring criteria to match when aggressiveness is set to Low.

For example, if the aggressiveness level is set to Low, each criterion for Low must be met the number of times

specified by the Recommendation Count value before automatically applying the recommendation.

If you set the Recommendation Count to **3**, Workload Balancing waits until all the criteria listed for Low are met

and repeated in three consecutive recommendations. This setting helps ensure that the VM actually needs to be moved

and that the recommended destination host has stable resource utilization over a longer period. It reduces the

potential for a recently moved VM to be moved off a host due to host performance changes after the move. By default,

this setting is set to Low to encourage stability.

We do not recommend increasing the **Optimization Aggressiveness** setting to increase the frequency with which your

hosts are being optimized. If you think that your hosts aren't being optimized quickly or frequently enough, try

adjusting the Critical thresholds. Compare the thresholds against the Pool Health report.

The consistency criteria associated with the different levels of aggressiveness is the following:

Low:

- All VMs in subsequent recommendations must be the same (as demonstrated by matching UUIDs in each recommendation).
- All destination hosts must be the same in subsequent recommendations
- The recommendation that immediately follows the initial recommendation must match or else the Recommendation Count reverts to 1

Medium:

- All VMs in subsequent recommendations must be from the same host; however, they can be different VMs from the ones in the first recommendation.
- All destination hosts must be the same in subsequent recommendations
- One of the next two recommendations that immediately follows the first recommendation must match or else the Recommendation Count reverts to 1

High:

- All VMs in the recommendations must be from the same host. However, the recommendations do not have to follow each other immediately.
- The host from which Workload Balancing recommended that the VM move must be the same in each recommendation
- The Recommendation Count remains at the same value even when the two recommendations that follow the first recommendation do not match

Optimization Aggressiveness example The following example illustrates how Workload Balancing uses the **Optimization Aggressiveness** setting and the Recommendation Count to determine whether to accept a recommendation automatically.

Each optimization recommendation issued by Workload Balancing proposes three VM placements. After these proposed placements, the recommendation count associated with each aggressiveness level is the number of times there have been consecutive recommendation at that Optimization Aggressiveness setting.

In the following examples, when the **Optimization Aggressiveness** setting is set to High, the Recommendation Count continues to increase after Recommendation 1, 2, and 3. This increase happens even though the same VMs are not recommended for new placements in each recommendation. Workload Balancing applies the placement recommendation with Recommendation 3 because it has seen the same behavior from that host for three consecutive recommendations.

In contrast, when set to Low aggressiveness, the consecutive recommendations count does not increase for the first four recommendations. The Recommendation Count resets to 1 with each recommendation because the same VMs were not recommended for placements. The Recommendation Count does not start to increase until the same recommendation is made in Recommendation #5. Finally, Workload Balancing automatically applies the recommendation made in Recommendation #6 after the third time it issues the same placement recommendations.

Recommendation 1:

Proposed placements:

- Move VM1 from host A to host B
- Move VM3 from host A to host B
- Move VM5 from host A to host C

Recommendation counts:

- High Aggressiveness Recommendation Count: 1
- Medium Aggressiveness Recommendation Count: 1
- Low Aggressiveness Recommendation Count: 1

Recommendation 2:

Proposed placements:

- Move VM1 from host A to host B
- Move VM3 from host A to host C
- Move VM7 from host A to host C

Recommendation counts:

- High Aggressiveness Recommendation Count: 2
- Medium Aggressiveness Recommendation Count: 1
- Low Aggressiveness Recommendation Count: 1

Recommendation 3:

Proposed placements:

- Move VM1 from host A to host B
- Move VM3 from host A to host C
- Move VM5 from host A to host C

Recommendation counts:

- High Aggressiveness Recommendation Count: 3 (Apply)
- Medium Aggressiveness Recommendation Count: 1
- Low Aggressiveness Recommendation Count: 1

Recommendation 4:

Proposed placements:

- Move VM1 from host A to host B
- Move VM3 from host A to host B
- Move VM5 from host A to host C

Recommendation counts:

- Medium Aggressiveness Recommendation Count: 2
- Low Aggressiveness Recommendation Count: 1

Recommendation 5:

Proposed placements:

- Move VM1 from host A to host B
- Move VM3 from host A to host B
- Move VM5 from host A to host C

Recommendation counts:

- Medium Aggressiveness Recommendation Count: 3 (Apply)
- Low Aggressiveness Recommendation Count: 2

Recommendation 6:

Proposed placements:

- Move VM1 from host A to host B
- Move VM3 from host A to host B
- Move VM5 from host A to host C

Recommendation counts:

- Low Aggressiveness Recommendation Count: 3 (Apply)

To configure VM recommendation intervals

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, click **Advanced**.
5. In the **VM Recommendation Interval** section, do one or more of the following:
 - In the **Minutes** box, type a value for the number of minutes Workload Balancing waits before making another optimization recommendation on a newly rebalanced host.
 - In the **Recommendations** box, type a value for the number of recommendations you want Workload Balancing to make before it applies a recommendation automatically.
 - Select a minimum severity level before optimizations are applied automatically.

- Modify how aggressively Workload Balancing applies optimization recommendations when it is running in automated mode. Increasing the aggressiveness level reduces constraints on the consistency of recommendations before automatically applying them. The **Optimization Aggressiveness** setting directly complements the **Recommendations** setting: that is, the recommendations count.

Note:

If you type “1” for the value in the **Recommendations** setting, the **Optimization Aggressiveness** setting is not relevant.

Adjust the Pool Audit Trail granularity settings

Follow this procedure to modify the granularity settings:

1. In XenCenter, select your pool.
2. In the **Properties** pane of the pool, click the **WLB** tab.
3. In the **WLB** tab, click **Settings**.
4. In the left pane, click **Advanced**.
5. On the **Advanced** page, click the **Pool Audit Trail Report Granularity** list, and select an option from the list.

Important:

Select the granularity based on your audit log requirements. For example, if you set your audit log report granularity to Minimum, the report only captures limited amount of data for specific users and object types. If you set the granularity to Medium, the report provides a user-friendly report of the audit log. If you choose to set the granularity to Maximum, the report contains detailed information about the audit log report. Setting the audit log report to Maximum can cause the Workload Balancing server to use more disk space and memory.

6. To confirm your changes, click **OK**.

View Pool Audit Trail reports based on objects in XenCenter

Follow this procedure to run and view reports of Pool Audit Trail based on the selected object:

1. After you have set the Pool Audit Trail Granularity setting, click **Reports**. The Workload Reports page appears.
2. Select **Pool Audit Trail** on the left pane.
3. You can run and view the reports based on a specific Object by choosing it from the **Object** list. For example, choose **Host** from the list to get the reports based on host alone.

Customize the event objects and actions captured by the Pool Audit Trail

To customize the event objects and actions captured by the Pool Audit Trail, you must sign in to the PostgreSQL database on the Workload Balancing virtual appliance, make the relevant changes to the list of event objects or actions, and then restart the Workload Balancing virtual appliance.

Sign in to the PostgreSQL database

1. Log on to the Workload Balancing virtual appliance console.
2. Run the following command:

```
1  psql -Upostgres -dWorkloadBalancing
```

3. Enter the database password. You set the database password when you ran the Workload Balancing configuration wizard after you imported the virtual appliance.

Customize event objects

Note:

In the command syntax that follows, `event_object` represents the name of the event object you want to add, update, or disable.

Enable an event object:

```
1  select * from update_audit_log_objects('event_object', true);
```

Disable an event object:

```
1 select * from update_audit_log_objects('event_object', false);
```

Get a list of event objects that are currently disabled:

```
1 select * from hv_audit_log_get_event_objects(false);
```

Get a list of event objects that are currently enabled:

```
1 select * from hv_audit_log_get_event_objects(true);
```

Customize event actions

Note:

In the command syntax that follows, `event_action` represents the name of the event action you want to add, update, or disable.

Enable an event action:

```
1 select * from update_audit_log_actions('event_action', true);
```

Disable an event action:

```
1 select * from update_audit_log_actions('event_action', false);
```

Get a list of event actions that are currently disabled:

```
1 select * from hv_audit_log_get_event_actions(false);
```

Get a list of event actions that are currently enabled:

```
1 select * from hv_audit_log_get_event_actions(true);
```

Restart the Workload Balancing virtual appliance Run the following commands to quit PostgreSQL and restart the Workload Balancing virtual appliance.

```
1 \q
```

```
1 systemctl restart workloadbalancing
```

Set alert level for Workload Balancing alerts in XenCenter

You can use the xe CLI to set the alert level for Workload Balancing alerts that are displayed in XenCenter.

Complete the following steps:

1. Run the following command on the pool coordinator to set the alert level for each alert code:

```
1 xe pool-send-wlb-configuration config:<wlb-alert-code>=<alert-level>
```

The 4 `wlb-alert-code` types are:

- MESSAGE_PRIORITY_WLB_OPTIMIZATION_ALERT - If Workload Balancing gives an optimization recommendation, this alert is raised.
- MESSAGE_PRIORITY_WLB_VM_RELOCATION - If Workload Balancing relocates a VM to another host, this alert is raised.
- MESSAGE_PRIORITY_WLB_HOST_POWER_OFF - If Workload Balancing optimization mode has been configured to `Maximize Density` and a host is powered off because there are no VMs running on the host, this alert is raised.
- MESSAGE_PRIORITY_WLB_HOST_POWER_ON - If Workload Balancing optimization mode has been configured to `Maximize Performance` and a host is powered on because doing so improves the host performance, this alert is raised.

The 6 `alert-level` types are:

- 0 - Mute the alert
- 1 - Critical
- 2 - Major
- 3 - Warning
- 4 - Minor
- 5 - Informational

2. Run the following command on the pool coordinator to view the alert levels set for the alert codes:

```
1 xe pool-retrieve-wlb-configuration
```

3. To test the alerts, raise a Workload Balancing alert and then, in XenCenter, click the `Notifications` panel to view the alert.

Administer Workload Balancing

June 26, 2026

After Workload Balancing has been running for a while, there are routine tasks that you might need to perform to keep

Workload Balancing running optimally. You might need to perform these tasks because of changes to your environment

(such as different IP addresses or credentials), hardware upgrades, or routine maintenance.

Connect to the Workload Balancing virtual appliance

After Workload Balancing configuration, connect the pool you want managed to the Workload Balancing virtual appliance

using either the CLI or XenCenter®. Likewise, you might need to reconnect to the same virtual appliance at some point.

To connect a pool to your Workload Balancing virtual appliance, you need the following information:

- IP address or FQDN of the Workload Balancing virtual appliance
 - To obtain the IP address for the Workload Balancing virtual appliance:
 1. In XenCenter, go to the Workload Balancing virtual appliance **Console** tab.
 2. Log in as `root` with the root password you created when you imported the appliance.
 3. Run the following command: `ifconfig`.
 - To specify the Workload Balancing FQDN when connecting to the Workload Balancing server, first add its host name and IP address to your DNS server.
- The port number of the Workload Balancing virtual appliance. By default, XenServer® connects to Workload Balancing on port 8012.

Only edit the port number when you have changed it during Workload Balancing Configuration.

The port number


specified during Workload Balancing Configuration, in any firewall rules, and in the Connect to WLB Server dialog

must match.

- Credentials for the resource pool you want Workload Balancing to monitor.

- Credentials for the Workload Balancing account you created during Workload Balancing configuration.

This account is often known as the Workload Balancing user account. XenServer uses this account to communicate with Workload Balancing. You created this account on the Workload Balancing virtual appliance during Workload Balancing Configuration.



When you first connect to Workload Balancing, it uses the default thresholds and settings for balancing workloads.

Automatic features, such as automated optimization mode, power management, and automation, are disabled by default.

Working with certificates

Certificate verification is enabled by default. To connect the Workload Balancing virtual appliance with XenServer, you must import the Workload Balancing certificate into your XenServer pool. If you are using XenCenter 2025.1 or later and Workload Balancing 8.4.0 or later, XenCenter automatically imports the required certificate from your Workload Balancing virtual appliance. If you are using earlier versions, you must import the certificate manually. For more information, see [Certificates](#).

If you want to upload a different (trusted) certificate or configure certificate verification, note the following before connecting your pool to Workload Balancing:

- The self-signed certificate generated during the first setup of the Workload Balancing virtual appliance is always valid for the IP address. If you configured the hostname and FQDN for the Workload Balancing virtual appliance during initial setup, the self-signed certificate is also valid for hostname and FQDN.

- If you want to use a certificate from a certificate authority, it is easier to specify the FQDN when connecting to Workload Balancing. However, you can specify a static IP address in the **Connect to WLB Server** dialog. Use this IP address as the Subject Alternative Name (SAN) in the certificate.

For more information, see [Certificates](#).

Disable certificate verification if your pool can't connect to Workload Balancing:

1. Get the pool UUID:

```
1 xe pool-list
```


2. Disable certificate verification:

```
1 xe pool-param-set wlb-verify-cert=false uuid=<uuid_of_pool>
```


If your pool can connect to Workload Balancing with certificate verification turned off, the issue is with your certificate configuration. If it cannot connect, the issue is with either your Workload Balancing credentials or your network connection. Contact Support or see [Troubleshoot Workload Balancing](#) for more information.

To connect your pool to the Workload Balancing virtual appliance

1. In XenCenter, select your resource pool and in its **Properties** pane, click the **WLB** tab. The **WLB** tab displays the **Connect** button.



2. In the **WLB** tab, click **Connect**. The **Connect to WLB Server** dialog box appears.



Connect to WLB Server

Server Address
Enter the address of the Workload Balancing server this Citrix Hypervisor resource pool will use.

Address:

Port: (Default is 8012)

WLB Server Credentials
Enter the credentials Citrix Hypervisor will use to connect to the Workload Balancing server.

Username:

Password:

Citrix Hypervisor Credentials
Enter the credentials the Workload Balancing Server will use to connect to Citrix Hypervisor.

Username:

Password:


☐ Use the current XenCenter credentials

OK Cancel

3. In the **Server Address** section, enter the following:

- a) In the **Address** box, type the IP address, hostname, or FQDN of the Workload Balancing virtual appliance.
For example, `WLB-appliance-computername.yourdomain.net`.
- b) (Optional) If you changed the Workload Balancing port during Workload Balancing Configuration, enter the port number in the **Port** box. XenServer uses this port to communicate with Workload Balancing.
By default, XenServer connects to Workload Balancing on port 8012.

4. In the **WLB Server Credentials** section, enter the user name and password that the pool uses to connect to the Workload Balancing virtual appliance.



☒ **Update Credentials**

WLB Server Credentials
Enter the credentials Citrix Hypervisor will use to connect to the Workload Balancing server.


Username:

Password:

These credentials must be for the account you created during Workload Balancing configuration. By default, the user name for this account is `wlbuser`.

5. In the **XenServer Credentials** section, enter the user name and password for the pool you are configuring.

Workload Balancing uses these credentials to connect to the hosts in that pool.



To use the credentials with which you are currently logged into XenServer, select **Use the current XenCenter credentials**. If you have assigned a role to this account using the role-based access control (RBAC) feature, ensure that the role has sufficient permissions to configure Workload Balancing. For more information, see [Workload Balancing Access Control Permissions](#).

6. Click **OK**. XenCenter starts the connection process.
7. Configure certificates:
 - (Workload Balancing 8.4.0 and later, XenCenter 2025.1.0 and later) If certificate verification is enabled for your pool, XenCenter now retrieves the Workload Balancing certificate from the virtual appliance.
 - a) If XenCenter displays the new or modified certificate, click **View Certificate** to review the information. You can verify the certificate fingerprint against the information in the Workload Balancing console.
 - b) Choose to **Accept** the certificate or **Cancel** the action.
 - If you choose to **Accept** the certificate, XenCenter installs it in your XenServer pool.
 - If you choose **Cancel**, XenCenter aborts the process. The certificate is not installed in the XenServer pool.

- (Earlier versions of Workload Balancing or XenCenter) Certificates are not automatically supported, you must configure XenServer to verify your Workload Balancing certificate. For more information, see [Certificates](#).

After connecting the pool to the Workload Balancing virtual appliance, Workload Balancing automatically begins monitoring the pool with the default optimization settings. If you want to modify these settings or change the priority given to resources, wait until the XenCenter Log shows that discovery is finished before proceeding.

Important:

After Workload Balancing is running for a time, if you do not receive optimal recommendations, evaluate your performance thresholds as described in [Configure Workload Balancing behavior](#). It is critical to set Workload Balancing to the correct thresholds for your environment or its recommendations might not be appropriate.

Workload Balancing access control permissions

When Role Based Access Control (RBAC) is implemented in your environment, all user roles can display the **WLB** tab.

However, not all roles can perform all operations. The following table lists the minimum role administrators require to use Workload Balancing features:

Permission	Minimum Required Role
Configure, Initialize, Enable, Disable WLB	Pool Operator
Apply WLB Optimization Recommendations in WLB tab	Pool Operator
Modify WLB report subscriptions	Pool Operator
Accept WLB Placement Recommendations	VM Power Admin
Generate WLB Reports, including the Pool Audit Trail report	Read Only
Display WLB Configuration	Read Only

The following table provides more details about permissions.

Permission	Allows Assignee To
Configure, Initialize, Enable, Disable WLB	Configure WLB Initialize WLB and change WLB servers Enable WLB Disable WLB
Apply WLB Optimization Recommendations in WLB tab	Apply any optimization recommendations that appear in the WLB tab
Modify WLB report subscriptions	Change the WLB report generated or its recipient
Accept WLB Placement Recommendations	Select one of the hosts Workload Balancing recommends for placement
Generate WLB Reports, including the Pool Audit Trail report	View and run WLB reports, including the Pool Audit Trail report
Display WLB Configuration	View WLB settings for a pool as shown on the WLB tab

If a user tries to use Workload Balancing and that user doesn't have sufficient permissions, a role elevation dialog appears. For more information about RBAC, see [Role-based access control](#).

Reconfigure a pool to use another Workload Balancing virtual appliance

You can reconfigure a resource pool to use a different Workload Balancing virtual appliance.

If you are moving from an older version of the Workload Balancing virtual appliance to the latest version, before disconnecting your old virtual appliance, you can migrate its data to the new version of the virtual appliance. For more information, see [Migrate data from an existing virtual appliance](#).

After disconnecting a pool from the old Workload Balancing virtual appliance, you can connect the pool by specifying the name of the new Workload Balancing virtual appliance.

To use a different Workload Balancing virtual appliance:

1. (Optional) Migrate data from an older version of the virtual appliance. For more information, see [Migrate data from an existing virtual appliance](#).
2. In XenCenter, from the **Pool** menu, select **Disconnect Workload Balancing Server** and click **Disconnect** when prompted.
3. In the **WLB** tab, click **Connect**. The **Connect to WLB Server** dialog appears.
4. Connect to the new virtual appliance. For more information, see [Connect to the Workload Balancing virtual appliance](#)

Change the Workload Balancing credentials

After initial configuration, if you want to update the credentials XenServer and the Workload Balancing appliance use to communicate, use the following process:

1. To pause Workload Balancing, go to the **WLB** tab and click **Pause**.
2. Change the Workload Balancing credentials by running the `wlbconfig` command. For more information, see [Workload Balancing Commands](#).
3. Re-enable Workload Balancing and specify the new credentials.
4. After the progress bar completes, click **Connect**.
The **Connect to WLB Server** dialog box appears.
5. Click **Update Credentials**.
6. In the **Server Address** section, modify the following settings as appropriate:
 - In the **Address** box, type the IP address or FQDN of the Workload Balancing appliance.
 - (Optional) If you changed the port number during Workload Balancing Configuration, enter that port number. The port number you specify in this box and during Workload Balancing Configuration is the port number XenServer uses to connect to Workload Balancing.

By default, XenServer connects to Workload Balancing on port 8012.

Note:

Only edit this port number if you changed it when you ran the Workload Balancing

Configuration wizard. The port number value specified when you ran the Workload Balancing Configuration wizard and the Connect to WLB Server dialog must match.

7. In the **WLB Server Credentials** section, enter the user name (for example, `wlbuser`) and password the computers running XenServer uses to connect to the Workload Balancing server.
8. In the **XenServer Credentials** section, enter the user name and password for the pool you are configuring (typically the password for the pool coordinator). Workload Balancing uses these credentials to connect to the hosts running XenServer in that pool.

To use the credentials with which you are currently logged into XenServer, select **Use the current XenCenter credentials**.

Change the Workload Balancing IP address

To change the Workload Balancing IP address, complete the following steps in the Workload Balancing virtual appliance console:

1. To view the current Workload Balancing IP address, run the `ifconfig` command on the virtual appliance.
2. Depending on the action you want to take run one of the following commands:

- To change DHCP to a static IP:

```
1 nmcli con mod eth0 ipv4.method manual ipv4.addresses <IP-address>/<CIDR-prefix> ipv4.gateway <Gateway>; nmcli con up eth0
```

- To change a static IP to DHCP:

```
1 nmcli con mod eth0 ipv4.method auto ipv4.addresses "" ipv4.gateway ""; nmcli con up eth0
```

- To change the static IP to another static IP:

```
1 nmcli con mod eth0 ipv4.method manual ipv4.addresses <new IP-address>/<CIDR-prefix> ipv4.gateway <Gateway>; nmcli con up eth0
```

3. To verify the change is complete, run the `ifconfig` command again to view the new Workload Balancing IP address.
4. To check that the Workload Balancing service is running normally, run the `systemctl status workloadbalancing` command.

If the returned result contains `Active: active (running)`, the Workload Balancing service is running normally. If the result contains `Active: inactive (dead)` or any other status, the Workload Balancing might exit abnormally.

Change the Workload Balancing virtual appliance configuration

When you first install the Workload Balancing virtual appliance it has the following default configuration:

Configuration	Value
Number of vCPUs	2
Memory (RAM)	2 GB
Disk space	30 GB

These values are suitable for most environments. If you are monitoring very large pools, you might consider increasing these values.

Change the number of vCPUs assigned to the virtual appliance

By default, the Workload Balancing virtual appliance is assigned 2 vCPUs. This value is sufficient for pools hosting 1000 VMs. You do not usually need to increase it. Only decrease the number of vCPUs assigned to the virtual appliance if you have a small environment.

This procedure explains how to change the number of vCPUs assigned to the Workload Balancing virtual appliance. Shut down the virtual appliance before performing these steps. Workload Balancing is unavailable for approximately five minutes.

1. Shut down the Workload Balancing virtual appliance.
2. In the XenCenter resource pane, select the Workload Balancing virtual appliance.
3. In the virtual appliance **General** tab, click **Properties**. The **Properties** dialog opens.
4. In the **CPU** tab of the **Properties** dialog, edit the CPU settings to the required values.
5. Click **OK**.
6. Start the Workload Balancing virtual appliance.

The new vCPU settings take effect when the virtual appliance starts.

Resize the virtual appliance memory

By default, the Workload Balancing virtual appliance is assigned 2 GB of memory.

For large pools, set the Workload Balancing virtual appliance to consume the maximum amount of memory you can make available to it (even up to 16 GB). Do not be concerned about high memory utilization. High memory utilization is normal for the virtual appliance because the database always consumes as much memory as it can obtain.

Note:

Dynamic Memory Control is not supported with the Workload Balancing virtual appliance. Set a fixed value for the maximum memory to assign to the virtual appliance.

This procedure explains how to resize the memory of the Workload Balancing virtual appliance. Shut down the virtual appliance before performing these steps. Workload Balancing is unavailable for approximately five minutes.

1. Shut down the Workload Balancing virtual appliance.
2. In the XenCenter resource pane, select the Workload Balancing virtual appliance.
3. In the virtual appliance **Memory** tab, click **Edit**. The **Memory Settings** dialog opens.
4. Edit the memory settings to the required values.
5. Click **OK**.
6. Start the Workload Balancing virtual appliance.

The new memory settings take effect when the virtual appliance starts.

Extend the virtual appliance disk

Warning:

You can only extend the available disk space in versions 8.3.0 and later as LVM is not supported before 8.3.0.

Workload Balancing does not support decreasing the available disk space.

By default, the Workload Balancing virtual appliance is assigned 30 GB of disk space.

The greater the number of VMs the Workload Balancing virtual appliance is monitoring, the more disk space it consumes per day.

You can estimate the amount of disk size that the virtual appliance needs by using the following formula:

```
1 Total estimated disk size = ( ( number of days * average disk usage ) +  
    base disk usage ) * grooming multiplier
```

- *number of days* is the number of days of data to retain
- *average disk usage* depends on the number of VMs being monitored. The following values give an approximation for certain numbers of VMs:
 - For 200 VMs - 0.246 GB/day
 - For 400 VMs - 0.505 GB/day
 - For 600 VMs - 0.724 GB/day
 - For 800 VMs - 0.887 GB/day
- *base disk usage* is 2.4 GB
- *grooming multiplier* is 1.25. This multiplier accounts for the amount of disk space required by grooming. It assumes that grooming requires an additional 25% of the total calculated disk space.

For versions 8.3.0 and later This procedure explains how to extend the virtual disk of the Workload Balancing virtual appliance for Workload Balancing versions 8.3.0 and later, using Linux Volume Manager (LVM).

Warning:

This procedure must only be followed by Experienced Linux System Administrators as incorrectly performing these steps can result in corrupting the Workload Balancing virtual appliance. Be sure to back up the

virtual appliance

before performing these steps and shut down the virtual appliance before you begin. Workload Balancing is unavailable for approximately five minutes.

To create new partitions, manipulate Physical Volumes and change your Filesystem size, perform the following actions

while logged in as a Super User (root):

1. View the current partitions:

```
1 fdisk -l
```

The default partitions might look like this:

```
-bash-4.2# fdisk -l

Disk /dev/xvda: 21.5 GB, 21474836480 bytes, 41943040 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000008b2

   Device Boot      Start         End      Blocks    Id  System
/dev/xvda1    *        2048     1026047       512000    83  Linux
/dev/xvda2          1026048     16777215       7875584    8e  Linux LVM

Disk /dev/mapper/centos-root: 7159 MB, 7159676928 bytes, 13983744 sectors
Units = sectors of 1 * 512 = 512 bytes
```

2. View the disk partition style:

```
1 parted <disk>
```

For example, to view the partition style of `/dev/xvda`:

```
1 parted /dev/xvda
```

3. Enter `p`.

If the following error messages occur, enter `Fix` to resolve each one:

- “Error: The backup GPT table is not at the end of the disk, as it should be. This might mean that another operating system believes the disk is smaller. Fix, by moving the backup to the end (and removing the old backup)?”
- “Warning: Not all of the space available to `<disk>` appears to be used, you can fix the GPT To use all of the

space (an extra <block number> blocks) or continue with the current setting?”

```
-bash-4.2# parted /dev/xvda
GNU Parted 3.1
Using /dev/xvda
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) print
Error: The backup GPT table is not at the end of the disk, as it should be. This might mean that another operating system believes the disk is smaller. Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel? Fix
Warning: Not all of the space available to /dev/xvda appears to be used, you can fix the GPT to use all of the space (an extra 20975616 blocks) or continue with the current setting?
Fix/Ignore? Fix
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 43.0GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name      Flags
 1      1049KB  3146KB  2098KB          grub      bios_grub
 2      3146KB  19.9MB  16.8MB    ext4        grubConfig
 3      19.9MB  32.2GB  32.2GB                roots

(parted) █
```

4. Enter **q** and press **Enter** to exit parted.

5. Edit the partitions:

```
1 fdisk <disk>
```

For example, to edit the partitions in Workload Balancing appliances:

```
1 fdisk /dev/xvda
```

6. Type **n** and press **Enter** to create a new partition, type **p** and press **Enter** to make it a primary partition, and press **Enter** to use the default which is the next available partition (in this case, as stated above, it'll be partition number 3).

Note:

If no additional space has been allocated yet, you will see a message indicating that there are no free sectors available. Type **q** and press **Enter** to quit fdisk. Allocate the desired space via XenCenter first and then come back to this step.

7. Press **Enter** twice to use the default first and last sectors of the available partition (or manually indicate the desired sectors). Type **t** to specify a partition type, choose the desired partition (in this case 3), type **8e**, and press **Enter** to make it an LVM type partition.

Example output:

```

Command (m for help): n
Partition type:
   p   primary (2 primary, 0 extended, 2 free)
   e   extended
Select (default p): p
Partition number (3,4, default 3):
First sector (16777216-41943039, default 16777216):
Using default value 16777216
Last sector, +sectors or +size{K,M,G} (16777216-41943039, default 41943039):
Using default value 41943039
Partition 3 of type Linux and of size 12 GiB is set

Command (m for help): t
Partition number (1-3, default 3):
Hex code (type L to list all codes): 8e
Changed type of partition 'Linux' to 'Linux LVM'

```

8. Type **p** and press **Enter** to print the details of the partition. The output should look similar to the one below (note that the start and end blocks values might vary depending on the amount of space you've allocated):

```

Command (m for help): p

Disk /dev/xvda: 21.5 GB, 21474836480 bytes, 41943040 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000008b2

   Device Boot      Start         End      Blocks    Id  System
/dev/xvda1    *        2048     1026047       512000    83   Linux
/dev/xvda2             1026048     16777215       7875584    8e   Linux LVM
/dev/xvda3       16777216     41943039     12582912    8e   Linux LVM

```

9. If something is incorrect, type **q** and press **Enter** to exit without saving and to prevent your existing partitions from being affected. Start again from step 1. Otherwise, if all looks well, type **w** and press **Enter** instead in order to write the changes.

After writing these changes, you might get a warning indicating that the device was busy and the kernel is still using the old table. If that's the case, run this command which will refresh the partition table, before proceeding with the next step: `partprobe`.

Make sure the new device partition (in this case `/dev/xvda4`) is listed now. To do so, run: `fdisk -l`.

The newly created device should be listed now:

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	1026047	512000	83	Linux
/dev/xvda2		1026048	16777215	7875584	8e	Linux LVM
/dev/ <u>xvda3</u>		16777216	41943039	12582912	8e	Linux LVM

10. If the output looks correct, create a Physical Volume:

```
1 pvcreate <new partition>
```

For example:

```
1 pvcreate /dev/xvda4
```

11. Check that the Physical Volume created above is now listed:

```
1 pvs
```

In this example, the additional space added was 12G. Example output:

```
-bash-4.2# pvcreate /dev/xvda3
Physical volume "/dev/xvda3" successfully created.
-bash-4.2# pvs
PV          VG      Fmt  Attr PSize  PFree
/dev/xvda2  centos  lvm2 a--  <7.51g 40.00m
/dev/xvda3          lvm2 ---  12.00g 12.00g
```

12. Based on the output of the previous command, the Volume Group named centos must be extended:

```
1 vgextend <volume group> <new partition>
```

For example:

```
1 vgextend centos /dev/xvda4
```

13. Check the current Volume Groups:

```
1 vgs
```

14. Run the following command:

```
1 pvscan
```

This should show `/dev/xvda4` as part of the centos Volume group. Example output:

```
-bash-4.2# vgextend centos /dev/xvda3
Volume group "centos" successfully extended
-bash-4.2# vgs
VG      #PV #LV #SN Attr   VSize  VFree
centos   2   2   0 wz--n- 19.50g <12.04g
-bash-4.2# pvscan
PV /dev/xvda2   VG centos      lvm2 [<7.51 GiB / 40.00 MiB free]
PV /dev/xvda3   VG centos      lvm2 [<12.00 GiB / <12.00 GiB free]
Total: 2 [19.50 GiB] / in use: 2 [19.50 GiB] / in no VG: 0 [0   ]
```

15. If the information shown in the previous steps looks correct, run this command to see the Logic Volume path for the Logical Volume to be extended:

```
1 lvdisplay
```

In this example, the path is `/dev/centos/root`:

```
--- Logical volume ---
LV Path                /dev/centos/root
LV Name                root
VG Name                centos
LV UUID                -
LV Write Access        read/write
LV Creation host, time localhost, 2017-04-26 11:
LV Status              available
# open                 1
LV Size                <6.67 GiB
Current LE             1707
Segments               1
Allocation              inherit
Read ahead sectors     auto
- currently set to     8192
Block device           253:0
```

16. Run the following command to view the free PE/size (this tells you the exact value to use when extending the partition):

```
1 vgdisplay
```

Example output:

```
Alloc PE / Size        1912 / <7.47 GiB
Free  PE / Size        3081 / <12.04 GiB
VG UUID                -----
```

17. Using the free PE/size value and the Logic Volume path outputted in step 11, extend the Logic Volume:

```
1 lvextend -l +100%FREE /dev/centos/root
```

If this executes successfully, extend the filesystem:

```
1 resize2fs /dev/centos/root
```

Example output:

```
-bash-4.2# lvextend -l +100%FREE /dev/centos/root
Size of logical volume centos/root changed from <30.00 GiB (7679 extents) to <40.02 GiB (10244 extents).
Logical volume centos/root successfully resized.
-bash-4.2# resize2fs /dev/centos/root
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/centos/root is mounted on /; on-line resizing required
old_desc_blocks = 4, new_desc_blocks = 6
The filesystem on /dev/centos/root is now 10489856 blocks long.
```

18. Verify that the filesystem size shows as expected:

```
1 df -h /*
```

```
-bash-4.2# df -h /*
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/centos-root    19G       4.1G    15G   22% /
/dev/xvda1                 497M      111M   387M   23% /boot
devtmpfs                  990M         0   990M    0% /dev
```

If you're seeing the expected numbers, you have successfully allocated the desired space and correctly extended the partition. For further assistance, please contact XenServer Support.

Stop Workload Balancing

Because Workload Balancing is configured at the pool level, when you want it to stop managing a pool, you must do one of the following:

Pause Workload Balancing. Pausing Workload Balancing stops XenCenter from displaying recommendations for the specified resource pool and managing the pool. Pausing is designed for a short period and lets you resume monitoring without having to reconfigure. When you pause Workload Balancing, data collection stops for that resource pool until you enable Workload Balancing again.

1. In XenCenter, select the resource pool for which you want to disable Workload Balancing.
2. In the **WLB** tab, click **Pause**. A message appears on the **WLB** tab indicating that Workload Balancing is

paused.

Tip:

To resume monitoring, click the **Resume** button in the **WLB** tab.

Disconnect the pool from Workload Balancing. Disconnecting from the Workload Balancing virtual appliance breaks

the connection between the pool and if possible, deletes the pool data from the Workload Balancing database. When you

disconnect from Workload Balancing, Workload Balancing stops collecting data on the pool.

1. In XenCenter, select the resource pool on which you want to stop Workload Balancing.
2. From the **Infrastructure** menu, select **Disconnect Workload Balancing Server**. The **Disconnect Workload Balancing server** dialog box appears.
3. Click **Disconnect** to stop Workload Balancing from monitoring the pool permanently.

Tip:

If you disconnected the pool from the Workload Balancing virtual appliance, to re-enable Workload Balancing on that

pool, you must reconnect to a Workload Balancing appliance. For information, see the [Connect to the Workload Balancing Virtual Appliance](#).

Enter maintenance mode with Workload Balancing enabled

With Workload Balancing enabled, if you put a host in maintenance mode, XenServer migrates the VMs running on that

host to their optimal hosts when available. XenServer uses Workload Balancing recommendations that are based on

performance data, your placement strategy, and performance thresholds to select the optimal host.

If an optimal host is not available, the words **Click here to suspend the VM** appear in the **Enter Maintenance**

Mode wizard. In this case, because there is not a host with sufficient resources to run the VM, Workload Balancing

does not recommend a placement. You can either suspend this VM or exit maintenance mode and suspend a VM on another

host in the same pool. Then, if you reenter the **Enter Maintenance Mode** dialog box, Workload Balancing might be

able to list a host that is a suitable candidate for migration.

Note:

When you take a host offline for maintenance and Workload Balancing is enabled, the words “Workload Balancing” appear in the **Enter Maintenance Mode** wizard.

To enter maintenance mode with Workload Balancing enabled:

1. In the **Resources** pane of XenCenter, select the physical server that you want to take off-line.
2. From the **Server** menu, select **Enter Maintenance Mode**.
3. In the **Enter Maintenance Mode** wizard, click **Enter maintenance mode**.

The VMs running on the host are automatically migrated to the optimal host based on the Workload Balancing performance data, your placement strategy, and performance thresholds.

To take the server out of maintenance mode:

1. Right-click the host and select **Exit Maintenance Mode**.

When you remove a host from maintenance mode, XenServer automatically restores that host's original VMs to that host.

Remove the Workload Balancing virtual appliance

To remove the Workload Balancing virtual appliance, we recommend that you use the standard procedure to delete VMs from XenCenter.

When you delete the Workload Balancing virtual appliance, the PostgreSQL database containing the Workload Balancing is deleted. To save this data, you must migrate it from the database before deleting the Workload Balancing virtual appliance.

Manage the Workload Balancing database

The Workload Balancing database is a PostgreSQL database. PostgreSQL is an open-source relational database. You can find documentation for PostgreSQL by searching the web.

The following information is intended for database administrators and advanced users of PostgreSQL who are comfortable with database administration tasks. If you are not experienced with PostgreSQL, we recommend that you become familiar with it before you attempt the database tasks in the sections that follow.

By default, the PostgreSQL user name is `postgres`. You set the password for this account during Workload Balancing configuration.

The amount of historical data you can store is based on the size of the virtual disk allocated to Workload Balancing and the minimum required space. By default, the size of the virtual disk allocated to Workload Balancing is 30 GB. In terms of managing the database, you can control the space that database data consumes by configuring database grooming. For more information, see [Database grooming parameters](#).

To store a lot of historical data, for example if you want to enable the Pool Audit trail Report, you can do either of the following:

- Make the virtual disk size assigned to the Workload Balancing virtual appliance larger. To do so, import the virtual appliance, and increase the size of the virtual disk by following the steps in [Extend the virtual appliance disk](#).
- Create periodic duplicate backup copies of the data by enabling remote client access to the database and using a third-party database administration tool.

Access the database

The Workload Balancing virtual appliance has a firewall (`firewalld`) configured in it. Before you can access the database, you must open the PostgreSQL server port (5432) in the appliance firewall.

1. From the Workload Balancing virtual appliance console, run the following command to open the PostgreSQL port:

```
1 firewall-cmd --permanent --zone=public --add-port=5432/tcp
```

2. Reload the firewall to apply the change:

```
1 firewall-cmd --reload
```

Control database grooming

The Workload Balancing database automatically deletes the oldest data whenever the virtual appliance reaches the minimum amount of disk space that Workload Balancing requires to run. By default, the minimum amount of required disk space is set to 4,096 MB.

The Workload Balancing database grooming options are controlled through the file `wlb.conf`.

When there is not enough disk space left on the Workload Balancing virtual appliance, Workload Balancing automatically starts grooming historical data. The process is as follows:

1. At a predefined grooming interval, the Workload Balancing data collector checks if grooming is required. Grooming is required if the database data has grown to the point where the only space that remains unused is the minimum required disk space. Use `GroomingRequiredMinimumDiskSizeInMB` to set the minimum required disk space.

You can change the grooming interval if desired using `GroomingIntervalInHour`. However, by default Workload Balancing checks to see if grooming is required once per hour.
2. If grooming is required, Workload Balancing begins by grooming the data from the oldest day. Workload Balancing then checks to see if there is now enough disk space for it to meet the minimum disk-space requirement.
3. If the first grooming did not free enough disk space, then Workload Balancing repeats grooming up to `GroomingRetryCounter` times without waiting for `GroomingIntervalInHour` hour.
4. If the first or repeated grooming freed enough disk space, then Workload Balancing waits for `GroomingIntervalInHour` hour and returns to Step 1.
5. If the grooming initiated by the `GroomingRetryCounter` did not free enough disk space, then Workload Balancing waits for `GroomingIntervalInHour` hour and returns to Step 1.

Database grooming parameters

There are five parameters in the `wlb.conf` file that control various aspects of database grooming. They are as follows:

- **GroomingIntervalInHour**. Controls how many hours elapse before the next grooming check is done. For example, if you enter **1**, Workload Balancing checks the disk space hourly. If you enter **2**, Workload Balancing checks disk space every two hours to determine if grooming must occur.
- **GroomingRetryCounter**. Controls the number of times Workload balancing tries rerunning the grooming database query.
- **GroomingDBDataTrimDays**. Controls the number of days worth of data Workload Balancing deletes from the database each time it tries to groom data. The default value is one day.
- **GroomingDBTimeoutInMinute**. Controls the number of minutes that the database grooming takes before it times out and is canceled. If the grooming query takes longer than is expected and does not finish running within the timeout period, the grooming task is canceled. The default value is 0 minutes, which means that database grooming never times out.
- **GroomingRequiredMinimumDiskSizeInMB**. Controls the minimum amount of free space left in the virtual disk assigned to the Workload Balancing virtual appliance. When the data in the virtual disk grows until there is only minimum disk size left on the virtual disk, Workload Balancing triggers database grooming. The default value is 4,096 MB.

To edit these values, see

[Edit the Workload Balancing configuration file.](#)

Change the database password

We recommend using the `wlbconfig` command to change the database password. For more information, see

[Modify the Workload Balancing configuration options.](#)

Do not change the password by modifying the `wlb.conf` file.

Archive database data

To avoid having older historical data deleted, you can, optionally, copy data from the database for archiving. To do so, you must perform the following tasks:

1. Enable client authentication on the database.
2. Set up archiving using the PostgreSQL database administration tool of your choice.

Enable client authentication to the database

While you can connect directly to the database through the Workload Balancing console, you can also use a PostgreSQL database management tool. After downloading a database management tool, install it on the system from which you want to connect to the database. For example, you can install the tool on the same laptop where you run XenCenter.

Before you can enable remote client authentication to the database, you must:

1. Modify the database configuration files, including `pg_hba.conf` file and the `postgresql.conf`, to allow connections.
2. Stop the Workload Balancing services, restart the database, and then restart the Workload Balancing services.
3. In the database management tool, configure the IP address of the database (that is, the IP address of the Workload Balancing virtual appliance) and the database password.

Modify the database configuration files

To enable client authentication on the database, you must modify the following files on the Workload Balancing virtual appliance: the `pg_hba.conf` file and the `postgresql.conf` file.

To edit the `pg_hba.conf` file:

1. Modify the `pg_hba.conf` file. From the Workload Balancing virtual appliance console, open the `pg_hba.conf` file with an editor, such as VI. For example:

```
1 vi /var/lib/pgsql/data/pg_hba.conf
```

2. If your network uses IPv4, add the IP address from the connecting computer to this file. For example:

In the configuration section, enter the following under `#IPv4 local connections`:

- **TYPE:** host
- **DATABASE:** all
- **USER:** all
- **CIDR-ADDRESS:** 0.0.0.0/0
- **METHOD:** trust

3. Enter your IP address in the `CIDR-ADDRESS` field.

Note:

Instead of entering 0.0.0.0/0, you can enter your IP address and replace the last three digits with 0/24. The trailing “24” after the / defines the subnet mask and only allows connections from IP addresses within that subnet mask.

When you enter `trust` for the `Method` field, it enables the connection to authenticate without requiring a password. If you enter `password` for the `Method` field, you must supply a password when connecting to the database.

4. If your network uses IPv6, add the IP address from the connecting computer to this file. For example:

Enter the following under `#IPv6 local connections`:

- **TYPE:** host
- **DATABASE:** all
- **USER:** all
- **CIDR-ADDRESS:** ::0/0
- **METHOD:** trust

Enter the IPv6 addresses in the `CIDR-ADDRESS` field. In this example, the `::0/0` opens the database up to connections from any IPv6 addresses.

5. Save the file and quit the editor.
6. After changing any database configurations, you must restart the database to apply the changes. Run the following command:

```
1 systemctl restart postgresql.service
```

To edit the `postgresql.conf` file:

1. Modify the `postgresql.conf` file. From the Workload Balancing virtual appliance console, open the `postgresql.conf` file with an editor, such as VI. For example:

```
1 vi /var/lib/pgsql/data/postgresql.conf
```

2. Edit the file so that it listens on any port and not just the local host. For example:
 - a) Find the following line:

```
1 # listen_addresses='localhost'
```

- b) Remove the comment symbol (#) and edit the line to read as follows:

```
1 listen_addresses='*'
```

3. Save the file and quit the editor.
4. After changing any database configurations, you must restart the database to apply the changes. Run the following command:

```
1 systemctl restart postgresql.service
```

Change the database maintenance window

Workload Balancing automatically performs routine database maintenance daily at 12:05AM GMT (00:05), by default.

During this maintenance window, data collection occurs but the recording of data might be delayed. However, the

Workload Balancing user interface controls are available during this period and Workload Balancing still makes optimization recommendations.

Note:

- During the maintenance window, the Workload Balancing server restarts. If you restart your VMs during this period, they will not be balanced by WLB until the WLB service has fully recovered.
- Similarly, at other times, if you restart all VMs in your pool, restarting the Workload Balanc-

ing server may
cause those VMs to remain unbalanced until WLB is available again.

Database maintenance includes releasing allocated unused disk space and reindexing the database. Maintenance lasts for approximately 6 to 8 minutes. In larger pools, maintenance might last longer, depending on how long Workload Balancing takes to perform discovery.

Depending on your time zone, you might want to change the time when maintenance occurs. For example, in the Japan Standard Time (JST) time zone, Workload Balancing maintenance occurs at 9:05 AM (09:05), which can conflict with peak usage in some organizations. If you want to specify a seasonal time change, such as Daylight Saving Time or summer time, you must build the change into value you enter.

To change the maintenance time:

1. In the Workload Balancing console, run the following command from any directory:

```
1 crontab -e
```

Workload Balancing displays the following:

```
1 05 0 * * * /opt/vpx/wlb/wlbmaintenance.sh
```

The value 05 0 represents the default time for Workload Balancing to perform maintenance in minutes (05) and then hours (0). (The asterisks represent the day, month, and year the job runs: Do not edit these fields.) The entry 05 0 indicates that database maintenance occurs at 12:05 AM, or 00:05, Greenwich Mean Time (GMT) every night. This setting means that if you live in New York, the maintenance runs at 7:05 PM (19:05) during winter months and 8:05 PM in summer months.

Important:

Do not edit the day, month, and year the job runs (as represented by asterisks). Database maintenance must run daily.

2. Enter the time at which you want maintenance to occur in GMT.
3. Save the file and quit the editor.

Customize Workload Balancing

Workload Balancing provides several methods of customization:

- **Command lines for scripting.** For more information, see [Workload Balancing commands](#).
- **Host Power On scripting support.** You can also customize Workload Balancing (indirectly) through the Host Power On scripting. For more information, see [Hosts and resource pools](#).

Upgrade Workload Balancing

Online upgrading of Workload Balancing has been deprecated for security reasons. Customers cannot upgrade by using the yum repo anymore. Customers can upgrade Workload Balancing to the latest version by importing the latest Workload Balancing virtual appliance downloadable at the [XenServer Downloads page](#).

Workload Balancing commands

This section provides a reference for the Workload Balancing commands. You can perform these commands from the XenServer host or console to control Workload Balancing or configure Workload Balancing settings on the XenServer host. This appendix includes xe commands and service commands.

Run the following service commands on the Workload Balancing appliance. To do so, you must log in to the Workload Balancing virtual appliance.

Log in to the Workload Balancing virtual appliance

Before you can run any service commands or edit the `wlb.conf` file, you must log in to the Workload Balancing virtual appliance. To do so, you must enter a user name and password. Unless you created extra user accounts on the virtual appliance, log in using the root user account. You specified this account when you ran Workload Balancing Configuration wizard (before you connected your pool to Workload Balancing). You can, optionally,

use the **Console** tab in XenCenter to log in to the appliance.

To log in to the Workload Balancing virtual appliance:

1. At the login prompt, enter the account user name.
2. At the Password prompt, enter the password for the account:

Note:

To log off from the Workload Balancing virtual appliance, simply type `logout` at the command prompt.

wlb restart

Run the `wlb restart` command from anywhere in the Workload Balancing appliance to stop and then restart the Workload Balancing Data Collection, Web Service, and Data Analysis services.

wlb start

Run the `wlb start` command from anywhere in the Workload Balancing appliance to start the Workload Balancing Data Collection, Web Service, and Data Analysis services.

wlb stop

Run the `wlb stop` command from anywhere in the Workload Balancing appliance to stop the Workload Balancing Data Collection, Web Service, and Data Analysis services.

wlb status

Run the `wlb status` command from anywhere in the Workload Balancing appliance to determine the status of the Workload Balancing server. After you run this command, the status of the three Workload Balancing services (the Web Service, Data Collection Service, and Data Analysis Service) is displayed.

Modify the Workload Balancing configuration options

Many Workload Balancing configurations, such as the database and web-service configuration options, are stored in the `wlb.conf` file. The `wlb.conf` file is a configuration file on the Workload Balancing virtual appliance.

To modify the most commonly used options, use the command `wlb config`. Running the `wlb config` command on the Workload Balancing virtual appliance lets you rename the Workload Balancing user account, change its password, or change the PostgreSQL password. After you run this command, the Workload Balancing services are restarted.

Run the following command on the Workload Balancing virtual appliance:

```
1 wlb config
```

The screen displays a series of questions guiding you through changing your Workload Balancing user name and password and the PostgreSQL password. Follow the questions on the screen to change these items.

Important:

Double-check any values you enter in the `wlb.conf` file: Workload Balancing does not validate values in the `wlb.conf` file. Therefore, if the configuration parameters you specify are not within the required range, Workload Balancing does not generate an error log.

Edit the Workload Balancing configuration file

You can modify Workload Balancing configuration options by editing the `wlb.conf` file, which is stored in `/opt/vpx/wlb` directory on the Workload Balancing virtual appliance. In general, only change the settings in this file with guidance from XenServer. However, there are three categories of settings you can change if desired:

- **Workload Balancing account name and password.** It is easier to modify these credentials by running the `wlb config` command.
- **Database password.** This value can be modified using the `wlb.conf` file. However, we recommend modifying it

through the `wlb config` command since this command modifies the `wlb.conf` file and automatically updates the password in the database. If you choose to modify the `wlb.conf` file instead, you must run a query to update the database with the new password.

- **Database grooming parameters.** You can modify database grooming parameters, such as the database grooming interval, using this file by following the instructions in the database management section. However, if you do so, we recommend using caution.

For all other settings in the `wlb.conf` file, we recommend leaving them at their default, unless you have been instructed to modify them.

To edit the `wlb.conf` file:

1. Run the following from the command prompt on the Workload Balancing virtual appliance (using VI as an example):

```
1 vi /opt/vpx/wlb/wlb.conf
```

The screen displays several different sections of configuration options.

2. Modify the configuration options, and exit the editor.

You do not need to restart Workload Balancing services after editing the `wlb.conf` file. The changes go into effect immediately after exiting the editor.

Important:

Double-check any values you enter in the `wlb.conf` file: Workload Balancing does not validate values in the `wlb.conf` file. Therefore, if the configuration parameters you specify are not within the required range, Workload Balancing does not generate an error log.

Increase the detail in the Workload Balancing log

The Workload Balancing log provides a list of events on the Workload Balancing virtual appliance, including actions for the analysis engine, database, and audit log. This log file is found in this location: `/var/log/wlb/LogFile.log`.

You can, if desired, increase the level of detail the Workload Balancing log provides. To do so, modify the

[Trace flags](#) section of the Workload Balancing configuration file ([wlb.conf](#)), which is found in the following

location: [/opt/vpx/wlb/wlb.conf](#). Enter a 1 or true to enable logging for a specific trace and a 0 or false to

disable logging. For example, to enable logging for the Analysis Engine trace, enter:

```
1 AnalEngTrace=1
```

You might want to increase logging detail before reporting an issue to XenServer Technical Support or when troubleshooting.

Logging Option	Trace Flag	Benefit or Purpose
Analysis Engine Trace	AnalEngTrace	Logs details of the analysis engine calculations. Shows details of the decisions the analysis engine is making and potentially gain insight into the reasons Workload Balancing is not making recommendations.
Database Trace	DatabaseTrace	Logs details about database reads/writes. However, leaving this trace on increases the log file size quickly.
Data Collection Trace	DataCollectionTrace	Logs the actions of retrieving metrics. This value lets you see the metrics Workload Balancing is retrieving and inserting into the Workload Balancing data store. However, leaving this trace on increases the log file size quickly.
Data Compaction Trace	DataCompactionTrace	Logs details about how many milliseconds it took to compact the metric data.

Logging Option	Trace Flag	Benefit or Purpose
Data Event Trace	<code>DataEventTrace</code>	This trace provides details about events Workload Balancing catches from XenServer.
Data Grooming Trace	<code>DataGroomingTrace</code>	This trace provides details about the database grooming.
Data Metrics Trace	<code>DataMetricsTrace</code>	Logs details about the parsing of metric data. Leaving this trace on increases the log-file size quickly.
Queue Management Trace	<code>QueueManagementTrace</code>	Logs details about data collection queue management processing. (This option is for internal use.)
Data Save Trace	<code>DataSaveTrace</code>	Logs details about the pool being saved to the database.
Score server Trace	<code>ScoreHostTrace</code>	Logs details about how Workload Balancing is arriving at a score for a host. This trace shows the detailed scores generated by Workload Balancing when it calculates the star ratings for selecting optimal hosts for VM placement.
Audit Log Trace	<code>AuditLogTrace</code>	Shows the action of the audit log data being captured and written. (This option is only for internal use and does not provide information that is captured in the audit log.) However, leaving this trace on increases the log file size quickly.

Logging Option	Trace Flag	Benefit or Purpose
Scheduled Task Trace	<code>ScheduledTaskTrace</code>	Logs details about scheduled tasks. For example, if your scheduled mode changes are not working, you might want to enable this trace to investigate the cause.
Web Service Trace	<code>WlbWebServiceTrace</code>	Logs details about the communication with the web-service interface.

Certificates for Workload Balancing

June 26, 2026

XenServer® and Workload Balancing communicate over HTTPS. During Workload Balancing Configuration, the wizard automatically creates a self-signed test certificate. This self-signed test certificate lets Workload Balancing establish a TLS connection to XenServer. By default, Workload Balancing creates this TLS connection with XenServer automatically. You do not need to perform any certificate configurations during or after configuration for Workload Balancing to create this TLS connection.

Note:

The self-signed certificate is a placeholder to facilitate HTTPS communication and is not from a trusted certificate authority. For added security, we recommend using a certificate signed from a trusted certificate authority.


To use a certificate from another certificate authority, such as a signed one from a commercial authority, you must configure Workload Balancing and XenServer to use it.

By default, XenServer validates the identity of the certificate before it establishes a connection to Workload

Balancing. To configure XenServer to check for a specific certificate, export the root certificate that was used to sign the certificate. Copy the certificate to XenServer and configure XenServer to check for it when a connection to Workload Balancing is made. XenServer acts as the client in this scenario and Workload Balancing acts as the server.

Depending on your security goals, you can either:

- [Configure XenServer to verify the self-signed certificate.](#)
- [Configure XenServer to verify a certificate-authority certificate.](#)



Note:

Certificate verification is a security measure designed to prevent unwanted connections. Workload Balancing certificates must meet strict requirements or the certificate verification doesn't succeed. When certificate verification fails, XenServer doesn't allow the connection.

For certificate verification to succeed, you must store the certificates in the specific locations in which XenServer expects to find the certificates.

Automatic certificate import

In XenCenter® 2025.1 and later with Workload Balancing 8.4.0 and later, XenCenter automatically imports the Workload Balancing certificate when first configuring the connection to Workload Balancing.

If the Windows system running XenCenter cannot access the Workload Balancing virtual appliance – for example, if the port used for communication between Workload Balancing and XenServer is disabled, automatic certificate import is not available and you must instead manually install the Workload Balancing certificate into your pool.

Configure XenServer to verify the self-signed certificate

You can configure XenServer to verify that the XenServer Workload Balancing self-signed certificate is authentic

before XenServer permits Workload Balancing to connect.

Important:

To verify the XenServer Workload Balancing self-signed certificate, you must connect to Workload Balancing using its host name. To find the Workload Balancing host name, run the `hostname` command on the virtual appliance.

To configure XenServer to verify the self-signed certificate, complete the following steps:

1. Copy the self-signed certificate from the Workload Balancing virtual appliance to the pool coordinator. The XenServer Workload Balancing self-signed certificate is stored at `/etc/ssl/certs/server.pem`. Run the following command on the pool coordinator:

```
1 scp root@<wlb-ip>:/etc/ssl/certs/server.pem wlb.pem
```

2. If you receive a message stating that the authenticity of `wlb-ip` cannot be established, type `yes` to continue.
3. Enter Workload Balancing virtual appliance root password when prompted. The certificate is copied to the current directory.
4. Install the certificate. Run the following command in the directory where you copied the certificate:

```
1 xe pool-certificate-install filename=wlb.pem
```

5. Verify the certificate was installed correctly by running the following command on the pool coordinator:

```
1 xe pool-certificate-list
```

If you installed the certificate correctly, the output of this command includes the exported root certificate.

Running this command lists all installed TLS certificates, including the certificate you installed.

6. To synchronize the certificate from the coordinator to all hosts in the pool, run the following command on the pool coordinator:

```
1 xe pool-certificate-sync
```

Running the `pool-certificate-sync` command on the coordinator synchronizes the certificate and certificate revocation lists on all the pool hosts with the coordinator. This action ensures all hosts in the pool use the same certificates.

There is no output from this command. However, the next step does not work if this one did not work successfully.

7. Instruct XenServer to verify the certificate before connecting to the Workload Balancing virtual appliance. Run the following command on the pool coordinator:

```
1 xe pool-param-set wlb-verify-cert=true uuid=uuid_of_pool
```

Tip:

Pressing the **Tab** key automatically populates the UUID of the pool.

8. (Optional) To verify this procedure worked successfully, perform the following steps:
 - a) To test if the certificate synchronized to the other hosts in the pool, run the `pool-certificate-list` command on those hosts.
 - b) To test if XenServer was set to verify the certificate, run the `pool-param-get` command with the `param-name=wlb-verify-cert` parameter. For example:

```
1 xe pool-param-get param-name=wlb-verify-cert uuid=uuid_of_pool
```

Configure XenServer to verify a certificate-authority certificate

You can configure XenServer to verify a certificate signed by a trusted certificate authority.

For trusted authority certificates, XenServer requires an exported certificate or certificate chain (the intermediate and root certificates) in `.pem` format that contains the public key.

If you want Workload Balancing to use a trusted authority certificate, do the following tasks:

1. Obtain a signed certificate from the certificate authority.
2. Specify and apply the new certificate.

3. Import the certificate chain into the pool.

Before beginning these tasks, ensure:

- You know the IP address for the XenServer pool coordinator.
- XenServer can resolve the Workload Balancing host name. (For example, you can try pinging the Workload Balancing FQDN from the XenServer console for the pool coordinator.)

Obtain a signed certificate from the certificate authority

To obtain a certificate from a certificate authority, you must generate a Certificate Signing Request (CSR). On the

Workload Balancing virtual appliance, create a private key and use that private key to generate the CSR.

Guidelines for specifying the Common Name The Common Name (CN) you specify when creating a CSR must exactly match the FQDN of your Workload Balancing virtual appliance. It must also match the FQDN or IP address you specified in the **Address** box in the **Connect to WLB Server** dialog box.

To ensure the name matches, specify the Common Name using one of these guidelines:

- Specify the same information for the certificate's Common Name as you specified in the **Connect to WLB Server** dialog.

For example, if your Workload Balancing virtual appliance is named `wlb-vpx.yourdomain`, specify `wlb-vpx.yourdomain` in the **Connect to WLB Server** dialog and provide `wlb-vpx.yourdomain` as the Common Name when creating the CSR.
- If you connected your pool to Workload Balancing by IP address, use the FQDN as the Common Name and the IP address as a Subject Alternative Name (SAN). However, this approach might not work in all situations.

Create a private key file On the Workload Balancing virtual appliance, complete the following steps:

1. Create a private key file:

```
1 openssl genrsa -des3 -out privatekey.pem 2048
```

2. Remove the password:

```
1 openssl rsa -in privatekey.pem -out privatekey.nop.pem
```

Note:

If you enter the password incorrectly or inconsistently, you might receive some messages indicating that there is a user interface error. You can ignore the message and rerun the command to create the private key file.

Generate the Certificate Signing Request On the Workload Balancing virtual appliance, complete the following steps:

1. Create the Certificate Signing Request (CSR) using the private key:

```
1 openssl req -new -key privatekey.nop.pem -out csr
```

2. Follow the prompts to provide the information necessary to generate the CSR:

Country Name. Enter the TLS Certificate country codes for your country. For example, CA for Canada or JM for Jamaica. You can find a list of TLS Certificate country codes on the web.

State or Province Name (full name). Enter the state or province where the pool is located. For example, Massachusetts or Alberta.

Locality Name. The name of the city where the pool is located.

Organization Name. The name of your company or organization.

Organizational Unit Name. Enter the department name. This field is optional.

Common Name. Enter the FQDN of your Workload Balancing server. This value must match the name the pool uses to connect to Workload Balancing. For more information, see [Guidelines for specifying the Common Name](#).

Email Address. This email address is included in the certificate when you generate it.

3. Provide optional attributes or click Enter to skip providing this information.

The CSR request is saved in the current directory and is named `csr`.

4. Display the CSR in the console window by running the following commands in the Workload Balancing appliance console:


```
1 cat csr
```

5. Copy the entire CSR and use it to request the certificate from the certificate authority.

Specify and apply the new certificate

Use this procedure to specify that Workload Balancing uses a certificate from a certificate authority. This procedure installs the root and (if available) intermediate certificates.

To specify a new certificate, complete the following steps:

1. Download the signed certificate, root certificate and, if the certificate authority has one, the intermediate certificate from the certificate authority.
2. If you didn't download the certificates directly to the Workload Balancing virtual appliance, copy them across by using one of the following methods:

- From a Windows computer, use WinSCP or another copying utility.

For the host name, you can enter the IP address and leave the port at the default. The user name and password are typically root and whatever password you set during configuration.

- From a Linux computer to the Workload Balancing appliance, use SCP or another copying utility. For example:

```
1 scp root_ca.pem root@wlb-ip:/path_on_your_WLB
```

3. On the Workload Balancing virtual appliance, merge the contents of all the certificates (root certificate, intermediate certificate - if it exists, and signed certificate) into one file. You can use the following command:

```
1 cat signed_cert.pem intermediate_ca.pem root_ca.pem > server.pem
```

4. Rename the existing certificate and key by using the move command:

```
1 mv /etc/ssl/certs/server.pem /etc/ssl/certs/server.pem_orig
2 mv /etc/ssl/certs/server.key /etc/ssl/certs/server.key_orig
```

5. Copy the merged certificate:

```
1 mv server.pem /etc/ssl/certs/server.pem
```

6. Copy the private key created previously:

```
1 mv privatekey.nop.pem /etc/ssl/certs/server.key
```

7. Make the private key readable only by root. Use the `chmod` command to fix permissions.

```
1 chmod 600 /etc/ssl/certs/server.key
```

8. Restart `stunnel`:

```
1 pkill -9 stunnel
2 stunnel
```

Import the certificate chain into the pool

After you obtain the certificates, import them onto the XenServer pool coordinator. Synchronize the hosts in the pool

to use those certificates. Then you can configure XenServer to check the certificate identity and validity each time

Workload Balancing connects to a host.

1. Copy the signed certificate, root certificate and, if the certificate authority has one, the intermediate certificate from the certificate authority onto the XenServer pool coordinator.

2. Install the root certificate on the pool coordinator:

```
1 xe pool-install-ca-certificate filename=root_ca.pem
```

3. If applicable, install the intermediate certificate on the pool coordinator:

```
1 xe pool-install-ca-certificate filename=intermediate_ca.pem
```

4. Verify both the certificates installed correctly by running this command on the pool coordinator:

```
1 xe pool-certificate-list
```

Running this command lists all installed TLS certificates. If the certificates installed successfully, they appear in this list.

5. Synchronize the certificate on the pool coordinator to all hosts in the pool:

```
1 xe pool-certificate-sync
```

Running the `pool-certificate-sync` command on the coordinator synchronizes the certificates and certificate

revocation lists on all the pool hosts with the pool coordinator. This action ensures all hosts in

the pool use
the same certificates.

6. Instruct XenServer to verify a certificate before connecting to the Workload Balancing virtual appliance. Run the following command on the pool coordinator:

```
1 xe pool-param-set wlb-verify-cert=true uuid=uuid_of_pool
```

Tip:

Pressing the Tab key automatically populates the UUID of the pool.

7. If you specified an IP address in the **Connect to WLB Server** dialog before you enabled certificate verification, you might be prompted to reconnect the pool to Workload Balancing.

Specify the FQDN for the Workload Balancing appliance in **Address** in the **Connect to WLB Server** dialog exactly as it appears in the certificate's Common Name. Enter the FQDN to ensure that the Common Name matches the name that XenServer uses to connect.

Troubleshooting

- Certificate verification is enabled by default on your pool. Disable certificate verification if your pool can't connect to Workload Balancing:

1. Get the pool UUID:

```
1 xe pool-list
```

2. Disable certificate verification:

```
1 xe pool-param-set wlb-verify-cert=false uuid=uuid_of_pool
```

If your pool can connect to Workload Balancing with certificate verification turned off, the issue is with your certificate configuration. If it cannot connect, the issue is with either your Workload Balancing credentials or your network connection. Contact Support or see [Troubleshoot Workload Balancing](#) for more information.

- Some commercial certificate authorities provide tools to verify the certificate installed correctly. Consider running these tools if these procedures fail to help isolate the issue. If these tools require specifying a TLS port, specify port 8012 or whatever port you set during Workload Balancing Configuration.
- If the **WLB** tab shows a connection error, there might be a conflict between the certificate Common Name and the name of the Workload Balancing virtual appliance. The Workload Balancing virtual appliance name and the Common Name of the certificate must match exactly.

For more information, see [Troubleshooting](#).

Regenerate the self-signed certificate

If the IP address of the Workload Balancing virtual appliance changes and no longer aligns with the value in the certificate, Workload Balancing shows a warning message in its console.

To address this warning, complete the following steps:

1. Log in to your Workload Balancing virtual appliance.
2. Regenerate the certificate by running the following command: `wlb_update_cert.sh`
3. Reconnect to the Workload Balancing virtual appliance.

Troubleshoot Workload Balancing

June 26, 2026

While Workload Balancing usually runs smoothly, this series of sections provides guidance in case you encounter issues.

Notes:

- Workload Balancing is available for XenServer® Premium Edition customers. For more information about XenServer licensing, see [Licensing](#). To upgrade, or to get a XenServer license, visit the [XenServer website](#).
- Workload Balancing 8.4.1 is compatible with XenServer 8.4 and XenServer 9.

Determine the status of the Workload Balancing virtual appliance

Run the `systemctl status workloadbalancing` command. For more information, see [Workload Balancing commands](#).

General troubleshooting tips

- Start troubleshooting by reviewing the Workload Balancing log files (`LogFile.log` and `wlb_install_log.log`). You can find these logs in the Workload Balancing virtual appliance in this location (by default):

`/var/log/wlb`

The level of detail in these log files can be configured by using the `wlb.conf` file. For more information, see

[Increase the detail in the Workload Balancing log](#).

- Check the logs in the XenCenter® **Notifications > Events** view for further information.
- To check the Workload Balancing virtual appliance build number, run the following command on a host in a pool that the virtual appliance monitors:

```
1 xe pool-retrieve-wlb-diagnostics | more
```

The Workload Balancing version number appears at the top of the output.

- The Workload Balancing virtual appliance is based on the CentOS operating system. If you experience CPU, memory, or disk related issues in the virtual appliance, you can use the standard Linux logs in `/var/log/*` to analyze the issue.
- Use standard Linux debugging and performance tuning commands to understand the virtual appliance behavior. For example, `top`, `ps`, `free`, `sar`, and `netstat`.

Error messages

Workload Balancing displays errors on screen as dialog boxes and as error messages in the **Notifications > Events** view in XenCenter.

If an error message appears, review the XenCenter event log for additional information. For more information, see the [XenCenter product documentation](#).

Issues entering Workload Balancing credentials

If you cannot successfully enter the virtual appliance user account and password while configuring the **Connect to**

WLB Server dialog, try the following:

- Ensure that the Workload Balancing virtual appliance is imported and configured correctly, and that all of its services are running.
- Check to ensure that you are entering the correct credentials. The **Connect to WLB Server** dialog asks for two different credentials:
 - **WLB Server Credentials:** XenServer uses this account to communicate with Workload Balancing. You created this account on the Workload Balancing virtual appliance during Workload Balancing Configuration. By default, the user name for this account is `wlbuser`.
 - **XenServer Credentials:** This account is used by the Workload Balancing virtual appliance to connect to the XenServer pool. This account is created on the XenServer pool coordinator and has the `pool-admin` or `pool-operator` role.
- You can enter the host name or FQDN of the Workload Balancing virtual appliance in the **Address** box, but this address must be one that the Windows system that runs XenCenter can resolve and access. Do not enter the host name of the physical server hosting the appliance. If you do not have a DNS entry for the FQDN, you must connect using the Workload Balancing appliance's IP address instead.
- Verify that the host is using the correct DNS server and the XenServer host can contact the Workload Balancing server using its FQDN. To do this check, ping the Workload Balancing appliance using its FQDN from the XenServer host. For example, enter the following in the XenServer host console:

```
1 ping wlb-vpx-1.mydomain.net
```

Issues with certificates

Certificate verification is enabled by default in XenServer 8.4 and later. To connect the Workload Balancing virtual appliance with XenServer, you must import the Workload Balancing certificate into your XenServer pool.

If your pool can't connect to Workload Balancing, verify whether your certificate configuration is the cause by disabling certificate checking for Workload Balancing:

1. Get the pool UUID:

```
1 xe pool-list
```

2. Disable certificate verification:

```
1 xe pool-param-set wlb-verify-cert=false uuid=<uuid_of_pool>
```

If your pool can connect to Workload Balancing with certificate verification turned off, the issue is with your certificate configuration. If it cannot connect, the issue is with either your Workload Balancing credentials or your network connection. Contact Support or see [Troubleshoot Workload Balancing](#) for more information.

Ensure that you enable certificate checking again after troubleshooting your issue:

1. Get the pool UUID:

```
1 xe pool-list
```

2. Enable certificate verification:

```
1 xe pool-param-set wlb-verify-cert=true uuid=<uuid_of_pool>
```

Issues with firewalls

The following error appears if the Workload Balancing virtual appliance is behind a hardware firewall, and you did not configure the appropriate firewall settings: “There was an error connecting to the Workload Balancing server: <pool

name> Click **Initialize WLB** to reinitialize the connection settings.”This error might also appear if the Workload Balancing appliance is otherwise unreachable.

If the Workload Balancing virtual appliance is behind a firewall, open port 8012.

Likewise, the port XenServer uses to contact Workload Balancing (8012 by default), must match the port number

specified when you ran the Workload Balancing Configuration wizard.

Workload Balancing connection errors


If you receive a connection error after configuring and connecting to Workload Balancing, the credentials might no

longer be valid. To isolate this issue:


1. Verify that the credentials you entered in the **Connect to WLB Server** dialog box are correct. For more information, see scenario 1 and 2.
2. Verify that the IP address or FQDN for the Workload Balancing virtual appliance that you entered in the **Connect to WLB Server** dialog box is correct.
3. Verify that the user name you created during Workload Balancing configuration matches the credentials you entered in the **Connect to WLB Server** dialog box.
4. If you receive a connection error in the Workload Balancing Status line on the **WLB** tab, you might need to reconfigure Workload Balancing on that pool. Click the **Connect** button on the **WLB** tab and reenter the host credentials.

You may encounter one of the following scenarios when attempting to establish a connection from XenCenter to the

Workload Balancing virtual appliance.

Scenario 1

This means that the credentials entered in the **XenServer Credentials** field in the **Connect to WLB Server** dialog box are incorrect. To fix this, double-check the credentials or check the **Use the current XenCenter credentials** box.

Scenario 2

This means that there is a problem with the credentials entered in the **WLB Server Credentials** field in the **Connect to WLB Server** dialog box when attempting to connect to the Workload Balancing virtual appliance (either the username or the password is incorrect). However, it can also mean that the Workload Balancing service is not running or that there is a problem with the database configuration file.

To fix credential issues, make sure that you are using the correct username and password. The default username for **WLB Server Credentials** field is `wlbuser` (not root). Root is the default administrator username. Note that

`wlbuser` is not an actual user with logon privileges in the appliance (it does not exist under `/etc/passwd`) and

thus these credentials are only used to connect to Workload Balancing itself. As such, they can be easily reset by

running the `wlbconfig` command. To change your credentials, see

[Change the Workload Balancing credentials](#).


To run the `wlbconfig` command, you must be able to log into the appliance as root. If the root password is unknown,

it can be reset using the regular CentOS/RHEL password recovery procedure.

If you have reset your credentials but the error still persists:

1. Check if the Workload Balancing process is running by using the `systemctl status workloadbalancing` command.
2. Make sure the `wlb.conf` file exists and is in the right directory by running this command:
`cat /opt/vpx/wlb/wlb.conf`

Scenario 3




This indicates that there is an issue connecting to the port specified under the Server Address options when

connecting to Workload Balancing from XenCenter (either the incorrect port was entered or the port is not listening).

To troubleshoot this:

1. Make sure the target appliance is up and running.
2. Double-check the port entered on the Workload Balancing connection details window (default is 8012).
3. Make sure this port is enabled in the appliance and listening. Use commands like `telnet <host> <port>` or `iptables -L` to help determine if the port is listening or if traffic is being denied on this port.

Scenario 4



This error occurs when there is a problem with stunnel (either it's not running or the certificate/key pair is incorrect). To troubleshoot this, first verify the certificate and key:

1. Confirm the certificate has not expired by running the following command:

```
1 openssl x509 -dates -in $(grep cert\ = /etc/stunnel/stunnel.conf |
  cut -d '=' -f2) -noout
```

2. Compare the hex on the output of the following 2 commands. If the output does not match then the wrong key is being used.

```
1 openssl x509 -modulus -in $(grep cert\ = /etc/stunnel/stunnel.conf
  | cut -d '=' -f2) -noout | openssl md5
```

and

```
1 openssl rsa -modulus -in $(grep key\ = /etc/stunnel/stunnel.conf |
  cut -d '=' -f2) -noout | openssl md5
```

If there are no problems with the certificate and key, make sure stunnel is running and is bound to port 8012 (or the configured port):


1. Run the following command in the WLB appliance CLI:

```
1 netstat -tulpn
```

On the output, 8012 (or the custom port) should show `status: LISTEN`.

2. If the appliance ran out of space, stunnel won't run. Use commands like `df -h` or `du -hs /*` to see whether you have enough space available on your appliance. To increase the disk space, see [Extend the virtual appliance disk](#).

Scenario 5



This error can occur because the stunnel process was terminated. If restarting the process yields the same results, restart the Workload Balancing virtual appliance.

Any other errors

If you encounter any other errors when attempting to connect to Workload Balancing or need further assistance

performing the steps above, collect the Workload Balancing logs which can be found under the `/var/log/wlb` directory in the Workload Balancing appliance.

Contact Support for further assistance.

Workload Balancing stops working

If Workload Balancing doesn't work (for example, it doesn't let you save changes to settings), check the Workload

Balancing log file for the following error message:

```
1 dwmdatcolsvc.exe: Don't have a valid pool. Trying again in 10 minutes.
```

This error typically occurs in pools that have one or more problematic VMs. When VMs are problematic, you might see the following behavior:

- **Windows.** The Windows VM crashes due to a stop error ("blue screen").
- **Linux.** The Linux VM might be unresponsive in the console and typically does not shut down.

To work around this issue:

1. Force the VM to shut down. To do so, you can do one of the following on the host with the problematic VM:

- In XenCenter, select the VM, and then from the VM menu, click **Force Shutdown**.
- Run the `vm-shutdown` xe command with the force parameter set to **true**. For example:

```
1  xe vm-shutdown force=true uuid=vm_uuid
```

You can find the host UUID on the **General** tab for that host (in XenCenter) or by running the `host-list`

xe command. You can find the VM UUID in the **General** tab for the VM or by running the `vm-list` xe command.

For more information, see [Command line interface](#).

2. In the `xsconsole` of the XenServer serving the crashed VM or in XenCenter, migrate all VMs to another host, then run the `xe-toolstack-restart` command. (Do not restart the toolstack while HA is enabled. If possible, temporarily disable HA before restarting the toolstack.)

Issues changing Workload Balancing servers

If you connect a pool to a different Workload Balancing server without disconnecting from Workload Balancing, both old and new Workload Balancing servers monitor the pool.

To solve this problem, you can take one of the following actions:

- Shut down and delete the old Workload Balancing virtual appliance.
- Manually stop the Workload Balancing services. These services are analysis, data collector, and Web service.

Note:

Do not use the `pool-deconfigure-wlb` xe command to disconnect a pool from the Workload Balancing virtual appliance or use the `pool-initialize-wlb` xe command to specify a different appliance.

Conversion Manager

June 25, 2026

Transform your VMware ESXi/vCenter VMs to XenServer® quickly and easily with the Conversion Manager virtual appliance.

Effortlessly batch-migrate multiple VMs in parallel, streamlining your entire migration process for maximum efficiency.

Learn more about moving from VMware to XenServer in our Tech Zone article:

[VMware to XenServer migration guide](#).

Conversion Manager enables you to:

- Batch convert multiple VMware ESXi/vCenter VMs in parallel using one simple wizard
- Map network settings between VMware and XenServer so your converted VMs can be up and running with the proper network settings
- Select a storage location where you would like your new XenServer VMs to run
- Optionally start each converted VM automatically after its conversion completes

Notes:

- XenCenter® does not remove or change your existing VMware environment. VMs are duplicated onto your XenServer environment and not removed from VMware. However, when converting Windows VMs with vTPM, the Conversion Manager removes the vTPM device from the original VMware Windows VM during the conversion process.
- During the migration, XenCenter prepares your VMs for networking and storage connectivity. After the conversion, the Conversion Manager automatically uninstalls VMware Tools from your VMs. To ensure full compatibility and optimal performance, we recommend installing XenServer VM Tools. These tools are vital for a fully supported setup, enhancing I/O performance, improving disk and network efficiency, and enabling key features such as clean shutdowns, reboots, suspensions, and live migrations. While VMs can run without XenServer VM Tools, performance might be limited. XenServer VM Tools are installed automatically on Linux VMs, and for Windows VMs, you can configure the Conversion Manager to install them automatically either from a version of the tools on your local disk or, if Conversion Manager has internet access, from the latest version of the tools available on the internet.
- Conversion Manager virtual appliance supports converting VMware ESXi/vCenter VMs with

different storage such as thin provisioning, thick provisioning, IDE, and SCSI.

- Conversion Manager virtual appliance does not require the source VMs to have VMware Tools installed. You can perform conversion on VMware ESXi/vCenter VMs regardless of whether they have VMware Tools installed.
- Conversion Manager virtual appliance cannot convert VMware ESXi/vCenter Windows VMs with four or more disks into XenServer VMs. Your VMware ESXi/vCenter Windows VMs must have three or fewer disks.
- After the conversion completes, the Conversion Manager appliance shuts down automatically to conserve host resources.

Supported guest operating systems

You can convert VMware ESXi/vCenter VMs running the following guest operating systems:

Guest operating system	Notes
Windows 10 (64-bit) Enterprise edition	Latest tested version is 22H2.
Windows 11 (64-bit) Enterprise edition	Only UEFI Boot mode or UEFI Secure Boot mode are supported.
Windows Server 2016 Standard (Desktop) edition	Only the listed Windows SKUs are supported for conversion.
Windows Server 2019 Standard (Desktop) edition	Only the listed Windows SKUs are supported for conversion.
Windows Server 2022 Standard (Desktop) edition	Only the listed Windows SKUs are supported for conversion.
Windows Server 2025 Standard (Desktop) edition	Preview feature. Only the listed Windows SKUs are supported for conversion.
Red Hat Enterprise Linux 8 (64-bit)	File system: EXT3, EXT4, or XFS Boot partition type: lvm or plain BIOS or UEFI boot mode is supported.
Red Hat Enterprise Linux 9 (64-bit)	File system: EXT3, EXT4, or XFS Boot partition type: lvm or plain BIOS or UEFI boot mode is supported.

Guest operating system	Notes
Rocky Linux 8 (64-bit)	File system: EXT3, EXT4, or XFS Boot partition type: lvm or plain BIOS or UEFI boot mode is supported.
Rocky Linux 9 (64-bit)	File system: EXT3, EXT4, or XFS Boot partition type: lvm or plain BIOS or UEFI boot mode is supported.
Ubuntu 22.04 (64-bit)	File system: EXT3, EXT4, or XFS Boot partition type: lvm or regular BIOS or UEFI boot mode is supported.
Ubuntu 24.04 (64-bit)	File system: EXT3, EXT4, or XFS Boot partition type: lvm or regular BIOS or UEFI boot mode is supported.

Note:

For Linux guests, UEFI Secure Boot is not supported for conversion. Only BIOS and UEFI boot modes are supported.

For more information about the guest operating systems supported by XenServer, see [Guest operating system support](#).

Understand XenServer

Before you can convert your environment, it is suggested that you become familiar with XenServer concepts. For more information, see [Technical overview](#).

To successfully convert VMware ESXi/vCenter VMs to XenServer, perform the following tasks:

- Set up a basic XenServer environment, including installing XenServer. For more information, see [Quick start](#) and [Install](#).
- Create a network in XenServer, assigning an IP address to a NIC. For more information, see [Quick start](#).
- Connect to storage. For more information, see [Quick start](#).

Compare VMware and XenServer terminology

The following table lists the approximate XenServer equivalent for common VMware features, concepts, and components:

VMware Term	XenServer Equivalent
VMware vSphere Client	XenCenter (the management console for XenServer)
VMware Tools	XenServer VM Tools (for Windows or Linux)
Cluster / Resource Pool	Resource Pool
Data Store	Storage Repository
vMotion	Live migration
Distributed Resource Scheduling (DRS)	Workload Balancing
High Availability (HA)	High Availability (HA)
vCenter Converter	Conversion Manager
Role Based Access Control (RBAC)	Role Based Access Control (RBAC)

Conversion overview

XenCenter and Conversion Manager virtual appliance create a copy of each targeted VM. After converting the targeted VM to a XenServer VM with comparable networking and storage connectivity, XenCenter imports the VM into your XenServer pool or host.

You can convert as few as one or two VMs or perform batch conversions of an entire environment. Simply select the VMs you want to convert and the Conversion Manager performs the conversions in batches of 10 VMs at a time, automatically selecting the next set of 10 VMs in the queue once the job is complete.

Note:

Before converting the VMs from vSphere, you must shut down the VMs (intended for conversion) on vSphere. Conversion Manager virtual appliance does not support converting a running VM using memory copied from vSphere to XenServer.

Also, before converting, ensure that a network and a storage controller exist in your VMware VM.


The conversion process requires four items:

- **XenCenter** - the XenServer management interface includes a conversion wizard where you set conversion options and control conversion. You can install XenCenter on your Windows desktop. XenCenter must be able to connect to XenServer and the Conversion Manager virtual appliance.
- **Conversion Manager virtual appliance** - a pre-packaged VM you import into the XenServer host or pool where you want to run the converted VMs. The virtual appliance converts the copies of the VMware ESXi/vCenter VMs into XenServer virtual machine format. After conversion, it imports these copies into the XenServer pool or host.
- **XenServer standalone host or pool** - the XenServer environment where you want to run the converted VMs.
- **VMware server** - Conversion Manager requires a connection to a VMware server that manages the VMs you want to convert. This connection can be to a vCenter Server, ESXi Server, or ESX Server. The VMs are not removed from the VMware server. Instead, the Conversion Manager Virtual Appliance makes a copy of these VMs and converts them to XenServer virtual-machine format.

Note:

When connecting to a vCenter Server, Conversion Manager accesses vCenter to retrieve the disk locations of the VMs on VMware, and then connects directly to the underlying vSphere (ESXi) host to read the disk data. Therefore, when using vCenter, you must configure Conversion Manager with network access to both the vCenter Server and all relevant vSphere hosts. Access to vCenter alone is not sufficient.

The following illustration shows the relationships between these components:



This illustration shows:

1. How XenCenter communicates with Conversion Manager virtual appliance.
2. How the Conversion Manager virtual appliance authenticates with the VMware server.
3. How the VMware server responds to the Conversion Manager virtual appliance during conversion.

The VMware server communicates with the Conversion Manager virtual appliance only when the appliance queries the VMware server for environment information and disk data throughout the conversion.

Summary of how to convert VMs

You can configure the Conversion Manager virtual appliance and start to convert VMs in just a few easy steps:

1. Download the latest version of the Conversion Manager virtual appliance from the [XenServer Downloads page](#).
2. Import the Conversion Manager virtual appliance into XenServer using XenCenter.
3. Configure the Conversion Manager virtual appliance by using XenCenter.
4. From XenCenter, launch the conversion wizard and start to convert VMs.

For more information on converting VMware ESXi/vCenter VMs, see [Get started with Conversion Manager](#).

What's new in Conversion Manager

June 29, 2026

Download the latest version of the Conversion Manager virtual appliance from the [XenServer Downloads page](#).

If you already have a previous version of the Conversion Manager virtual appliance installed and want to upgrade to the latest version, there is no automatic upgrade path. Download the latest version of the virtual appliance and remove the older version from your system.

For more information on converting VMware ESXi/vCenter VMs, see [Get started with Conversion Manager](#).

What's new in 2026.1.0

Released Jul 6, 2026

This release includes the following new features:

- Support for converting additional Linux guests: Red Hat Enterprise Linux 9, Rocky Linux 8 and 9, and Ubuntu 22.04 and 24.04.
- Support for converting Linux guests that use the XFS file system or UEFI boot mode.
- Support for converting VMs from VMware 9 to XenServer®.
- Support for converting VMs to XenServer 9.
- Support for converting a single virtual disk of up to 16 TiB on GFS2 storage repositories.
- A new option in XenCenter® to automatically start each VM after its conversion completes.
- (Preview) Support for converting Windows Server 2025 guests.

To use these features, ensure that your XenCenter is updated to the latest version.

Get started with Conversion Manager

June 25, 2026

You can easily convert your VMware ESXi/vCenter virtual machines (VMs) to XenServer® in just a few steps:

1. [Prepare your XenServer environment and review the prerequisite information.](#)
2. [Import and configure the Conversion Manager virtual appliance by using XenCenter.](#)

Note:

If you already have a previous version of the Conversion Manager virtual appliance installed and want to upgrade to the latest version, there is no automatic upgrade path. Download the latest version of the virtual appliance from the [XenServer downloads page](#) and remove the older version from your system.

3. [From XenCenter, launch the conversion wizard and begin converting your VMware ESXi/vCenter VMs to XenServer.](#)
4. [Complete the post-conversion tasks.](#)
5. [Review other conversion tasks.](#)

Prepare your environment

Before converting your VMware environment, you must create and prepare the target XenServer stand-alone host or pool to run the converted VMware ESXi/vCenter VMs. Preparing your environment includes the following activities:

1. Defining a strategy of how you convert your VMware environment. Do you want to convert 1 or 2 VMs? Do you want to convert your entire environment? Do you want to create a pilot first to ensure that your configuration is correct? Do you run both environments in parallel? Do you want to maintain your existing cluster design when you convert to XenServer?
2. Planning your networking configuration. Do you want to connect to the same physical networks? Do you want to simplify or change your networking configuration?
3. Installing XenServer on the hosts you want in the pool. Ideally, plug the NICs on the hosts into their physical networks before you begin installation.

4. Creating a pool and performing any basic networking configuration. For example, do the following:
 - Configure a network to connect to the VMware cluster on the XenServer host (if the cluster is not on the same network as the XenServer host).
 - Configure a network to connect to the storage array. That is, if you use IP-based storage, create a XenServer network that connects to the physical network of the storage array.
 - Create a pool and add hosts to this pool.
5. (For shared storage and XenServer pools.) Preparing the shared storage where you store the virtual disks and creating a connection to the storage, known as a Storage Repository (SR) on the pool.
6. (Optional) Although not a requirement for conversion, you might want to configure the administrator accounts on the XenServer pool to match those accounts on the VMware server. For information about configuring Role-based Access Control for Active Directory accounts, see [Role-based access control](#).

Install XenServer and create a pool

Before you can convert VMware ESXi/vCenter VMs, ensure that you create a XenServer pool or host where you want to run the converted VMs. This pool must have networking configured so it can connect to the VMware server. You might also want to configure the same physical networks on the XenServer pool that you have in the VMware cluster, or simplify your networking configuration. If you want to run the converted VMs in a pool, create a storage repository before conversion and add the shared storage to the pool.

If you are new to XenServer, you can learn about XenServer basics, including basic installation and configuration, by reading [Quick start](#).

XenServer environment considerations

Before installing XenServer and importing the virtual appliance, consider the following factors that might change your

conversion strategy:

Guest operating systems supported for conversion You can convert VMware ESXi/vCenter VMs running the following guest operating systems:

- Windows 10 (64-bit) Enterprise edition
- Windows 11 (Only UEFI Boot mode or UEFI Secure Boot mode are supported.)
- Windows Server 2016 Standard (Desktop) edition
- Windows Server 2019 Standard (Desktop) edition
- Windows Server 2022 Standard (Desktop) edition
- Windows Server 2025 Standard (Desktop) edition (Preview)
- Red Hat Enterprise Linux 8 (64-bit) with the following configuration:
 - File system: EXT3, EXT4, or XFS
 - Boot partition type: lvm or plain
 - BIOS or UEFI boot mode
- Red Hat Enterprise Linux 9 (64-bit) with the following configuration:
 - File system: EXT3, EXT4, or XFS
 - Boot partition type: lvm or plain
 - BIOS or UEFI boot mode
- Rocky Linux 8 (64-bit) with the following configuration:
 - File system: EXT3, EXT4, or XFS
 - Boot partition type: lvm or plain
 - BIOS or UEFI boot mode
- Rocky Linux 9 (64-bit) with the following configuration:
 - File system: EXT3, EXT4, or XFS
 - Boot partition type: lvm or plain
 - BIOS or UEFI boot mode
- Ubuntu 22.04 (64-bit) with the following configuration:
 - File system: EXT3, EXT4, or XFS
 - Boot partition type: lvm or regular
 - BIOS or UEFI boot mode
- Ubuntu 24.04 (64-bit) with the following configuration:
 - File system: EXT3, EXT4, or XFS
 - Boot partition type: lvm or regular
 - BIOS or UEFI boot mode

Note:

- For Linux guests, UEFI Secure Boot is not supported for conversion. Only BIOS and UEFI boot modes are supported.
- VMware does not distinguish between Ubuntu or Rocky Linux versions. When you convert an Ubuntu or Rocky Linux VM, select the correct version manually before conversion.

For more information about the guest operating systems supported by XenServer, see [Guest operating system support](#).

Host requirements Select the host where you want to run the Conversion Manager virtual appliance. This host can be a stand-alone host or a host in the pool where you run the converted VMs.

For pools, you can run the virtual appliance on any host in the pool, provided its storage meets the storage requirements.

Note:

We recommend that you run only one Conversion Manager in a pool at a time.

The Conversion Manager virtual appliance requires a minimum of:

- XenServer 8.4
- Disk space: 30 GB of disk space
- Memory: 6 GB
- Virtual CPU allocation: 2 vCPU

Networking requirements To convert VMware ESXi/vCenter VMs, the Conversion Manager virtual appliance needs connectivity to a physical network or VLAN that can contact the VMware server. (In the following sections, this network is referred to as the “VMware network”.)

If the VMware server is on a different physical network than the hosts in the XenServer pool, add the network to XenServer before conversion.

Note:

- The time it takes for your VMs to be converted depends on the physical distance between your VMware and XenServer networks and also the size of your VM's virtual disk. You can estimate how long the conversion will last by testing the network throughput between your VMware server and XenServer.
- By default, the Conversion Manager uses HTTPS to download the VM's virtual disk during VM conversion. To speed up the migration process, you can switch the download path to HTTP.

Map your existing network configuration Conversion Manager virtual appliance includes features that can reduce the amount of manual networking configuration needed after you convert from your existing VMware ESXi/vCenter VMs to XenServer. For example, the Conversion Manager virtual appliance will:

- Preserve virtual MAC addresses on the VMware ESXi/vCenter VMs and reuse them in the resulting XenServer VMs.
Preserving the MAC addresses associated with virtual network adapters (virtual MAC addresses) may:

- Help preserve IP addresses in environments using DHCP
- Be useful for software programs whose licensing references the virtual MAC addresses

- Map (virtual) network adapters. Conversion Manager virtual appliance can map VMware networks onto XenServer networks so that after the VMs are converted, their virtual network interfaces are connected accordingly.

For example, if you map VMware 'Virtual Network 4' to XenServer 'Network 0', any VMware VM that had a virtual adapter connected to 'Virtual Network 4' is connected to 'Network 0' after conversion. Conversion Manager virtual appliance does not convert or migrate any hypervisor network settings. The wizard only alters a converted VM's virtual network interface connections based on the mappings provided.

Note:

You do not need to map all of your VMware networks on to the corresponding XenServer networks. However, if you

prefer, you can change the networks the VMs use, reduce, or consolidate the number of networks in your new XenServer configuration.

To gain the maximum benefit from these features, we recommend the following:

- Before installing XenServer, plug the hosts into the networks on the switch (that is, the ports) that you would like to configure on the host.
- Ensure that the XenServer pool can see the networks that you would like to be detected. Specifically, plug the XenServer hosts into switch ports that can access the same networks as the VMware cluster.

Though it is easier to plug the XenServer NICs into the same networks as the NICs on the VMware hosts, it is not required. If you would like to change the NIC/network association, you can plug a XenServer NIC into a different physical network.

Prepare for the Conversion Manager virtual appliance networking requirements When you perform a conversion, you must create a network connection to the network where the VMware server resides.

Conversion Manager virtual appliance uses this connection for conversion traffic between the XenServer host and the VMware server.

To create this network connection, you must perform two tasks:

- When you import the Conversion Manager virtual appliance, specify the network you added for conversion traffic as a virtual network interface. You can do so by configuring **interface 1** so it connects to that network.
- Before you run the conversion wizard, add the network connecting VMware and XenServer to the XenServer host where you want to run the converted VMs.

By default, when you import the Conversion Manager virtual appliance, XenCenter® creates one virtual network interface associated with Network 0 and NIC0. When adding a network for conversion, select a network other than XenServer's

management network to improve performance in busy pools. For more information about the management interface, see

[Networking](#).

Inside the Conversion Manager, you might see multiple network interfaces (eth0 and eth1). eth0 attaches to the host's

internal network that is used to communicate with the local dom0. eth1 attaches to the routable network that is used to

communicate with XenCenter.

To add a network to XenServer:

1. In the **Resource** pane in XenCenter, select the pool where you would like to run Conversion Manager virtual appliance.
2. Click the **Networking** tab.
3. Click **Add Network**.
4. On the **Select Type** page, select **External Network**, and click **Next**.
5. On the **Name** page, enter a meaningful name for the network (for example, "VMware network") and a description.
6. On the **Interface** page, specify the following:
 - **NIC**. The NIC that you want XenServer to use to create the network. Select the NIC that is plugged in to the physical or logical network of the VMware server.
 - **VLAN**. If the VMware network is a VLAN, enter the VLAN ID (or "tag").
 - **MTU**. If the VMware network uses jumbo frames, enter a value for the Maximum Transmission Unit (MTU) between 1500 and 9216. Otherwise, leave the MTU box at its default value of 1500.

Note:

Do not select the **Automatically add this network to new virtual machines** checkbox.

7. Click **Finish**.

Storage requirements Before you convert batches of VMware ESXi/vCenter VMs, consider your storage requirements.

- If the converted VMs run on a single standalone host (not a pool), you can specify either local or shared storage as the location for the converted virtual disks.
- If you want to run the converted VMs in a pool, their virtual disks must be stored on shared storage. Ensure that you add the shared storage to the pool by creating a storage repository.

This storage repository must be large enough to contain the virtual disks for all the converted VMs you want to run in that pool.

To create a storage repository:

1. In the **Resource** pane in XenCenter, select the pool where you intend to run the Conversion Manager virtual appliance.
2. Click the **Storage** tab.
3. Click **New SR** and follow the instructions in the wizard. For more instructions, press **F1** to display the online help.

XenServer requirements You can run VMs converted with this release of Conversion Manager on the following versions of XenServer:

- XenServer 8.4
- XenServer 9

VMware requirements Conversion Manager virtual appliance can convert VMware ESXi/vCenter VMs from the following versions of VMware:

- vCenter Server 7.x, 8.x, and 9.x
- vSphere 7.x, 8.x, and 9.x
- ESXi 7.x, 8.x, and 9.x

Note:

Conversion Manager virtual appliance cannot convert VMware ESXi/vCenter Windows VMs with four or more disks into XenServer VMs. Your VMware ESXi/vCenter Windows VMs must have three or fewer disks. Your VMware ESXi/vCenter VMs must also have a network and a storage controller configured.

Import and configure the virtual appliance

The Conversion Manager virtual appliance is a single pre-installed VM designed to run on a XenServer host.

Download the virtual appliance

The Conversion Manager virtual appliance is packaged in XVA format. You can download the virtual appliance from the [XenServer downloads page](#). When downloading the file, save it to a folder on your local hard drive (typically, but not necessarily, on the computer where XenCenter is installed). After the `.xva` file is on your hard drive, you can import it into XenCenter.

Note:

Conversion Manager virtual appliance is available for XenServer Premium Edition customers or those customers who have access to XenServer through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about XenServer licensing, see [Licensing](#). To upgrade, or to get a XenServer license, visit the Citrix website.

Import the virtual appliance into XenServer

To import the Conversion Manager virtual appliance into the pool or host where you want to run the converted VMs, use the XenCenter **Import** wizard:

1. Open XenCenter. Right-click on the pool (or host) into which you want to import the virtual appliance package, and select **Import**.
2. Browse to locate the virtual appliance package.
3. Select the pool or a *home server* where you want to run the Conversion Manager virtual appliance.

Note:

A home server is the host that provides the resources for a VM in a pool. While it can, a XenServer attempts to start the VM on that host, before trying other hosts. If you select a host, the Conversion

Manager virtual

appliance uses this host as its home server. If you select the pool, the virtual appliance automatically starts on the most suitable host in that pool.

4. Choose a storage repository on which to store the virtual disk for the Conversion Manager virtual appliance and then click **Import**. To add a storage repository to the pool, see the section called “Meet Storage Requirements.” You can choose either local or shared storage.
5. Ensure the network to be used for conversion (which connects the VMware server to the XenServer host) is selected as the network associated with **interface 1** (“virtual NIC 1”).
 - If the correct network does not appear beside interface 1, use the list in the **Network** column to select a different network.
 - If you have not added the VMware network that is on a different physical network than the pool, do the following:
 - a) Exit the wizard.
 - b) Add the network to the pool.
 - c) Rerun the wizard.

For more information, see **To add a network to XenServer**.

Warning:

Do NOT configure NIC0 to your customer network. Assign NIC0 only to “Host internal management network.”

6. Leave the **Start VM after import** checkbox enabled, and click **Finish** to import the virtual appliance.
7. After importing the `.xva` file, the Conversion Manager virtual appliance appears in the **Resources** pane in XenCenter.

Configure the Conversion Manager virtual appliance

Before you can use the Conversion Manager virtual appliance to convert VMware ESXi/vCenter VMs, configure it using the XenCenter **Console** tab:

1. After importing the Conversion Manager virtual appliance, click the **Console** tab.
2. Read the license agreement. To view the contents of the license agreement, open the URL in a web browser. Press any key to continue.
3. Enter and confirm a new root password for the Conversion Manager virtual appliance. We recommend selecting a strong password.
4. Enter a host name for the Conversion Manager virtual appliance.
5. Enter the domain suffix for the virtual appliance. For example, if the fully qualified domain name (FQDN) for the virtual appliance is `citrix-migrate-vm.domain4.example.com`, enter `domain4.example.com`.
6. Enter **y** to use DHCP to obtain the IP address automatically for the Conversion Manager virtual appliance. Otherwise, enter **n** and then enter a static IP address, subnet mask, and gateway for the VM.
7. Review the host name and network setting and enter **y** when prompted. This step completes the Conversion Manager virtual appliance configuration process.
8. When you have successfully configured the appliance, a login prompt appears. Enter the login credentials and press **Enter** to log in to the Conversion Manager virtual appliance.

If you do not configure the Conversion Manager correctly, you might have to reconfigure the virtual appliance. For more information, see [Reconfigure the Conversion Manager virtual appliance](#).

Reconfigure the Conversion Manager virtual appliance If you need to change the configuration of the Conversion Manager after installing the virtual appliance, you can follow these steps to reconfigure it:

1. Log in to the Conversion Manager virtual appliance by using SSH or the XenCenter **Console** tab as the root user.
2. Delete the following file: `/etc/init.d/vpx_startup_setup_done`.
3. Reboot the virtual appliance by using XenCenter.
4. Reconfigure the virtual appliance by using the XenCenter **Console** tab.

Convert VMware ESXi/vCenter VMs

When you convert VMware ESXi/vCenter VMs, they are imported into the XenServer pool or standalone host where you are running the Conversion Manager virtual appliance. Converted VMs retain their original VMware settings for the virtual processor and virtual memory.

Before you start the conversion procedure, ensure that the following requirements are met:

- You have credentials for the XenServer pool or standalone host. Either the root account or an RBAC (Role-Based Access Control) account with the Pool Admin role is acceptable.
- You have credentials for the VMware server hosting the VMs to be converted. These are needed to connect the Conversion Manager Console to the VMware server.
- The VMware virtual machines to be converted:
 - Are powered off.
 - Have both a network and a storage controller configured.
 - For Windows VMs:
 - Log in to the VMware virtual machine as an administrator.
 - Disable the “Use Admin Approval Mode for the built-in Administrator Account” group policy.
- The XenServer pool (or host) that runs the converted VMs is connected to a storage repository. The storage repository must contain enough space for the converted virtual disks.
- If you want to run your newly converted VMs in a pool, the storage repository must be shared storage. However, if the converted VMs run on a single standalone host (not a pool), you can use local storage.
- Each virtual disk of the VM to convert is within the supported maximum size. On a GFS2 storage repository, a single virtual disk can be up to 16 TiB. On other storage repository types, each virtual disk must be less than 2 TiB.
- The XenServer pool (or host) has networks that the converted VMs use.

To convert your VMware ESXi/vCenter VMs into VMs that can run in a XenServer environment:

1. If the VMware VM is encrypted, decrypt it before starting the conversion. If the VM has a vTPM attached, remove the vTPM before decrypting the VM.

2. Ensure that the virtual appliance is installed and running on the XenServer host or pool where you want to import the VMs.

3. In XenCenter, go to **Pool > Conversion Manager**.

The **Conversion Manager** window opens. Wait while the wizard connects to your virtual appliance.

4. Click **New Conversion**.

5. In the **New Conversion** wizard, enter the credentials for the VMware server:

- **Server.** Enter the IP address or FQDN for the VMware server that contains the VMs you want to convert to XenServer.
- **Username.** Enter a valid user name for this VMware server. This account must either be a VMware admin account or have a Root role.
- **Password.** Enter the password for the user account you specified in the **Username** box.

Click **Next**. XenCenter connects to the VMware server.

6. On the **Virtual Machines** page, select the VMs you want to convert from the list available on the VMware server.

The conversion fails if the VM name contains unsupported characters such as %, &, *, \$, #, @, !, \, /, :, *, ?, ", <, >, |, ;, or '. Only uppercase and lowercase ASCII characters, numbers, spaces, periods (.), hyphens (-), underscores (_), plus signs (+), and parentheses (()) are allowed.

Note:

If BitLocker is enabled on your VM, or if you are converting a VM with an attached vTPM used by BitLocker, we recommend shutting down BitLocker before conversion. This action decrypts the VM disk and enables the Conversion Manager to automatically uninstall VMware Tools and install XenServer VM Tools on your Windows VMs.

Click **Next**.

7. If you are converting a Windows VM with an attached vTPM, review the mandatory tasks and information on the **vTPM** page. (This page only appears when converting a Windows VM with an attached vTPM.)

To proceed with the conversion, you must connect to the vCenter Server. vCenter access is necessary for the

Conversion Manager to remove the existing vTPM and create a XenServer vTPM for the VM.

Note:

If other applications rely on the vTPM, back up your vTPM data before conversion, as the Conversion Manager cannot recover encrypted data from the old vTPM. A new empty vTPM is created on XenServer. After the conversion, manually restore your data to the new vTPM on XenServer.

Once you have completed all the mandatory tasks, check the agreement checkbox at the bottom of the page to confirm

your completion and then click **Next**.

8. On the **Template** page, select the template to use for each of the selected VMs. By default, the recommended template for the VM is selected. For more information, see Template mappings.

For Windows VMs, this page also lets you choose whether to receive updates to the I/O drivers (a component of the XenServer VM Tools) through Windows Update. By default, this option is selected. For Citrix Provisioning targets or reset-on-boot machines, we recommend that you deselect this option.

For Linux VMs, this checkbox is disabled.

9. On the **Storage** page, select the storage repository you want to use during conversion. This storage repository is where the VMs and the virtual disks that you are creating are stored permanently.

This tab indicates the proportion of available storage that the virtual disks of the converted VMs consume.

10. On the **Networking** page, for each VMware network listed, select the XenServer network to map it to. You can also select whether to preserve virtual MAC addresses. Click **Next**.

11. If you are converting Windows VMs, enable the **Install XenServer VM Tools for Windows after conversion** checkbox on the **XenServer VM Tools** page. This allows the Conversion Manager to automatically install XenServer VM Tools on your Windows VMs. (This page only appears for Windows VM conversions.)

Installing XenServer VM Tools for Windows on each Windows VM is essential for a fully supported configuration. These

tools enhance I/O performance, improving disk and network efficiency while supporting key VM features such as clean shutdowns, reboots, suspensions, and live migrations. While VMs can operate without these tools, their performance might be limited. If you enable automatic installation, choose a source:

- **Use the existing XenServer VM Tools version** - If you have previously downloaded and uploaded XenServer VM Tools for Windows, you can select that version. The version you uploaded is displayed for your reference.
- **Download and upload the latest version of XenServer VM Tools for Windows** - If internet access is available, XenCenter automatically connects to the XenServer website to download and upload the latest version of the tools to the Conversion Manager.
- **Select a tools file from disk** - When you need a specific version or lack internet connectivity, you can manually upload the tools file from a local disk.

Click **Next**.

12. On the **Finish** page, review the options you configured for the conversion process. You can click **Previous** to change these options.

(Optional) Select **Start the converted VMs automatically after conversion** to power on each VM as soon as its conversion completes successfully. This option requires Conversion Manager 2026.1.0 or later and is disabled by default.

Note:

If a conversion job fails or does not complete, the converted VM does not start automatically, even when this option is enabled.

To proceed with the configuration shown, click **Finish**.

The conversion process begins, during which the Conversion Manager automatically uninstalls VMware Tools on your Windows or Linux VMs. When the conversion completes, the converted VMs start automatically if you selected the auto-start option. Otherwise, start your VMs manually from XenCenter by right-clicking the VM

and selecting

Start.

After the VMs boot, XenServer VM Tools install automatically to ensure compatibility. To finalize the installation, manually restart the VM once more.

The **Conversion Manager** window displays conversions in progress and completed conversions.

Post-conversion tasks

After converting your VMs, complete the following tasks on Windows and Linux VMs:

Reactivate the Windows License on Windows VMs

On Windows VMs, depending on your Microsoft licensing model, you might have to reactivate the VM's Windows license. This reactivation happens because the Windows operating system perceives the conversion as a hardware change.

Reboot your Windows VMs

For Windows VMs, if you chose to have the Conversion Manager automatically install XenServer VM Tools, the tools are installed after the initial boot. To finalize the installation, manually restart the VM once more.

Other conversion tasks

The **Manage Conversions** window enables you to perform other tasks related to converting VMs. These tasks include clearing jobs, saving a summary of jobs, retrying jobs, canceling jobs, and displaying the log file.

To clear all jobs:

1. Select **Clear All**.
2. When prompted to confirm this action, click **Yes** to continue.

To save a summary of jobs:

1. Click **Export All**.
2. Specify where to save the CSV file.
3. Click **Save**.

To retry a job:

- 1. Select the job from the list.
- 2. Click **Retry**.

Note:

The **Retry** option is only enabled for failed or canceled jobs.

To cancel a job:

- 1. Select the job from the list.
- 2. Click **Cancel**.

Note:

Cancel jobs is only enabled for queued or running jobs.

To save the conversion log file for a single job:

- 1. Select the job from the list.
- 2. From the logs menu, Click **Fetch Selected Log**.
- 3. Specify where to save the log file.

To save the conversion log file for all jobs:

- 1. From the logs menu, Click **Fetch All Logs**.
- 2. Specify where to save the log file.

To display conversion details:

- 1. Select the job from the list.
- The information is displayed in the **Details** panel.

Template mappings

The following table shows the mapping between the template of your VMware VM and the suggested XenServer template:

VMware template name	XenServer template name
Microsoft Windows Server 2025 (64-bit)	Windows Server 2025 (64-bit)
Microsoft Windows Server 2022 (64-bit)	Windows Server 2022 (64-bit)
Microsoft Windows Server 2019 (64-bit)	Windows Server 2019 (64-bit)

VMware template name	XenServer template name
Microsoft Windows Server 2016 (64-bit)	Windows Server 2016 (64-bit)
Microsoft Windows 11 (64-bit)	Windows 11
Microsoft Windows 10 (64-bit)	Windows 10 (64-bit)
Red Hat Enterprise Linux 9 (64-bit)	Red Hat Enterprise Linux 9
Red Hat Enterprise Linux 8 (64-bit)	Red Hat Enterprise Linux 8
Rocky Linux (64-bit)	Rocky Linux 8 or Rocky Linux 9
Ubuntu Linux (64-bit)	Ubuntu Jammy Jellyfish 22.04 or Ubuntu Noble Numbat 24.04

Notes

1. If you attempt to convert a VM that is not supported by the Conversion Manager, the **New Conversion** wizard suggests “Other install media” as the template.
2. The VMware Ubuntu and Rocky Linux templates do not specify a version, so the suggested template might not be the most appropriate. Select the correct version manually before conversion.

Troubleshoot Conversion Manager

June 25, 2026

This section provides information about troubleshooting the conversion process and converted VMs.

Problems starting a converted VM

In general, conversion runs smoothly and Conversion Manager virtual appliance converts VMs without any issues. However, in some rare cases, you might receive errors when attempting to open converted VMs. The following sections provide some guidance on resolving errors and other issues.

Blue screen with Windows STOP code 0x0000007B

This stop code indicates that Conversion Manager virtual appliance was unable to configure a Windows device that is critical to boot in XenServer® for the first time. Save the logs and send them to Support for further guidance.

Windows product activation

Depending on your licensing model, an error message on system activation might appear when you attempt to start a Windows VM.

Lost network settings in a Windows VM

If you import a Windows VM from an ESXi server to XenServer, the IPv4/IPv6 network settings can be lost. To retain the network settings, reconfigure the IPv4/IPv6 settings after completing the conversion.

Unable to start VMware SCSI disk

If a VMware VM boots from a SCSI disk but also has IDE hard disks configured, the VM might not boot when you convert it to XenServer. This issue occurs because the migration process assigns the IDE hard disks lower device numbers than SCSI disks. However, XenServer boots from the hard disk assigned to device 0. To resolve this issue, rearrange the virtual disk position in XenCenter® so that the VM reboots from the virtual disk that contains the operating system.

To change the position of the virtual disk containing the operating system:

1. In the XenCenter **Resources** pane, select the powered off guest VM.
2. Select the **Storage** tab.
3. From the **Virtual Disks** list, select the virtual disk containing the operating system and then click **Properties**.
4. In the virtual disk's **Properties** dialog, click the **vm_name** tab to display device options.
5. From the **Device Position** list, select **0** and Click **OK**.

Problems during conversion

If you experience problems or errors when converting VMs, try exporting the VMware VM as an OVF package. If you cannot export the VMware VM as an OVF package, Conversion Manager cannot convert this VM. Use the error messages you receive when attempting to export the VM as an OVF package to troubleshoot and fix the issues with your VMware VM. For example, you might have to configure a network or a storage controller before the VM can be exported as an OVF package or converted. For more information about troubleshooting your VMware ESXi/vCenter VMs, see the [VMware documentation](#).

If you see any errors when converting Linux VMs, remove the converted VM, restart the Conversion Manager virtual appliance and retry.

Logs of failed conversions are stored in the Conversion Manager virtual appliance and can be retrieved by clicking

Fetch All Logs on the **Conversion Manager** window. When you contact Support to raise any issues, we recommend that you provide the conversion log file and, additionally, a full server status report for troubleshooting. For more information, see [Creating a Server Status Report](#).

Command-line interface

June 22, 2026

The xe CLI enables you to script and automate system administration tasks. Use the CLI to integrate XenServer® into an existing IT infrastructure.

Getting started with the xe CLI

The xe command line interface is installed by default on all XenServer hosts. A remote Windows version is included with XenCenter. A stand-alone remote CLI is also available for Linux.

On your XenServer host

The `xe` command line interface is installed by default on your host. You can run `xe` CLI commands in the dom0 console. Access the dom0 console in one of the following ways:

- In XenCenter, go to the **Console** tab for the host where you want to run the command.
- SSH into the host where you want to run the command.

On Windows

On Windows, the `xe.exe` command is installed along with XenCenter.

To use the `xe.exe` command, open a Windows Command Prompt and change directories to the directory where the `xe.exe` file is located (typically `C:\Program Files (x86)\XenServer\XenCenter`). If you add the `xe.exe` installation location to your system path, you can use the command without having to change into the directory.

On Linux

On RPM-based distributions (such as Red Hat), you can install the stand-alone `xe` command from the RPM named `client_install/xapi-xe-BUILD.x86_64.rpm` on the main XenServer installation ISO.

To install from the RPM, use the following command:

```
1 rpm -ivh xapi-xe-BUILD.x86_64.rpm
```

You can use parameters at the command line to define the XenServer host, user name, and password to use when running `xe` commands. However, you also have the option to set this information as an environment variable. For example:

```
1 export XE_EXTRA_ARGS="server=<host name>,username=<user name>,password=<password>"
```

Note:

The remote `xe` CLI on Linux might hang when attempting to run commands over a secure connection and these commands involve file transfer. If so, you can use the `--no-ssl` parameter to run the command over an insecure connection to the XenServer host.

Getting help with `xe` commands

Basic help is available for CLI commands on-host by typing:

```
1 xe help command
```

A list of the most commonly used xe commands is displayed if you type:

```
1 xe help
```

Or a list of all xe commands is displayed if you type:

```
1 xe help --all
```

Basic xe syntax

The basic syntax of all XenServer xe CLI commands is:

```
1 xe command-name argument=value argument=value
```

Each specific command contains its own set of arguments that are of the form `argument=value`. Some commands have required arguments, and most have some set of optional arguments. Typically a command assumes default values for some of the optional arguments when invoked without them.

If the xe command runs remotely, extra arguments are used to connect and authenticate. These arguments also take the form `argument=argument_value`.

The `server` argument is used to specify the host name or IP address. The `username` and `password` arguments are used to specify credentials.

A `password-file` argument can be specified instead of the password directly. In this case, the xe command attempts to read the password from the specified file and uses that password to connect. (Any trailing CRs and LFs at the end of the file are stripped off.) This method is more secure than specifying the password directly at the command line.

The optional `port` argument can be used to specify the agent port on the remote XenServer host (defaults to 443).

Example: On the local XenServer host:

```
1 xe vm-list
```

Example: On a remote XenServer host:

```
1 xe vm-list username=username password=password server=hostname
```

Shorthand syntax is also available for remote connection arguments:

- `-u` user name
- `-pw` password

- `-pwf` password file
- `-p` port
- `-s` server

Example: On a remote XenServer host:

```
1 xe vm-list -u myuser -pw mypassword -s hostname
```

Arguments are also taken from the environment variable `XE_EXTRA_ARGS`, in the form of comma-separated key-value pairs. For example, to enter commands that are run on a remote XenServer host, first run the following command:

```
1 export XE_EXTRA_ARGS="server=jeffbeck,port=443,username=root,password=pass"
```

After running this command, you no longer have to specify the remote XenServer host parameters in each `xe` command that you run.

Using the `XE_EXTRA_ARGS` environment variable also enables tab completion of `xe` commands when issued against a remote XenServer host, which is disabled by default.

Special characters and syntax

To specify argument/value pairs on the `xe` command line, write: `argument=value`

Unless the value includes spaces, do not use quotes. Do not include whitespace between the argument name, the equals sign (=), and the value. Any argument not conforming to this format is ignored.

For values containing spaces, write: `argument="value with spaces"`

When you use the CLI on your XenServer host, commands have a tab completion feature similar to the feature in the standard Linux bash shell. For example, if you type `xe vm-l` and then press the **TAB** key, the rest of the command is displayed. If more than one command begins with `vm-l`, pressing **TAB** a second time lists the possibilities. This feature is useful when specifying object UUIDs in commands.

Note:

Tab completion does not normally work when running commands on a remote XenServer host. However, if you set the `XE_EXTRA_ARGS` variable on the machine where you enter the commands, tab completion is enabled. For more information, see [Basic xe syntax](#).

Command types

The CLI commands can be split in two halves. Low-level commands are concerned with listing and parameter manipulation of API objects. Higher level commands are used to interact with VMs or hosts in a more abstract level.

The low-level commands are:

- *class-list*
- *class-param-get*
- *class-param-set*
- *class-param-list*
- *class-param-add*
- *class-param-remove*
- *class-param-clear*

Where *class* is one of:

- *bond*
- *console*
- *host*
- *host-crashdump*
- *host-cpu*
- *network*
- *patch*
- *pbk*
- *pif*
- *pool*
- *sm*
- *sr*
- *task*
- *template*
- *vbd*
- *vdi*

- `vif`
- `vlan`
- `vm`

Not every value of *class* has the full set of *class-param-action* commands. Some values of *class* have a smaller set of commands.

Parameter types

The objects that are addressed with the `xe` commands have sets of parameters that identify them and define their states.

Most parameters take a single value. For example, the `name-label` parameter of a VM contains a single string value. In the output from parameter list commands, such as `xe vm-param-list`, a value in parentheses indicates whether parameters are read-write (RW) or read-only (RO).

The output of `xe vm-param-list` on a specified VM might have the following lines:

```
1 user-version ( RW): 1
2 is-control-domain ( RO): false
```

The first parameter, `user-version`, is writable and has the value 1. The second, `is-control-domain`, is read-only and has a value of false.

The two other types of parameters are multi-valued. A *set* parameter contains a list of values. A *map* parameter is a set of key-value pairs. As an example, look at the following piece of sample output of the `xe vm-param-list` on a specified VM:

```
1 platform (MRW): acpi: true; apic: true; pae: true; nx: false
2 allowed-operations (SRO): pause; clean_shutdown; clean_reboot; \
3 hard_shutdown; hard_reboot; suspend
```

The `platform` parameter has a list of items that represent key-value pairs. The key names are followed by a colon character (:). Each key-value pair is separated from the next by a semicolon character (;). The M preceding the RW indicates that this parameter is a map parameter and is readable and writable. The `allowed-operations` parameter has a list that makes up a set of items. The S preceding the RO indicates that this is a set parameter and is readable but not writable.

To filter on a map parameter or set a map parameter, use a colon (:) to separate the map parameter name and the key-value pair. For example, to set the value of the `foo` key of the `other-config` parameter of a VM to `baa`, the command would be

```
1 xe vm-param-set uuid=VM uuid other-config:foo=baa
```

Low-level parameter commands

There are several commands for operating on parameters of objects: *class-param-get*, *class-param-set*, *class-param-add*, *class-param-remove*, *class-param-clear*, and *class-param-list*. Each of these commands takes a **uuid** parameter to specify the particular object. Since these commands are considered low-level commands, they must use the UUID and not the VM name label.

- **xe class-param-list** *uuid=uuid*

Lists all of the parameters and their associated values. Unlike the *class-list* command, this command lists the values of “expensive” fields.

- **xe class-param-get** *uuid=uuid param-name=parameter param-key=key*

Returns the value of a particular parameter. For a map parameter, specifying the *param-key* gets the value associated with that key in the map. If *param-key* is not specified or if the parameter is a set, the command returns a string representation of the set or map.

- **xe class-param-set** *uuid=uuid param=value*

Sets the value of one or more parameters.

- **xe class-param-add** *uuid=uuid param-name=parameter key=value param-key=key*

Adds to either a map or a set parameter. For a map parameter, add key-value pairs by using the *key=value* syntax. If the parameter is a set, add keys with the *param-key=key* syntax.

- **xe class-param-remove** *uuid=uuid param-name=parameter param-key=key*

Removes either a key-value pair from a map, or a key from a set.

- **xe class-param-clear** *uuid=uuid param-name=parameter*

Completely clears a set or a map.

Low-level list commands

The *class-list* command lists the objects of type *class*. By default, this type of command lists all objects, printing a subset of the parameters. This behavior can be modified in the following ways:

- It can filter the objects so that it only outputs a subset
- The parameters that are printed can be modified.

To change the parameters that are printed, specify the argument *params* as a comma-separated list of the required parameters. For example:

```
1 xe vm-list params=name-label,other-config
```

Alternatively, to list all of the parameters, use the syntax:

```
1 xe vm-list params=all
```

The list command doesn't show some parameters that are expensive to calculate. These parameters are shown as, for example:

```
1 allowed-VBD-devices (SR0): <expensive field>
```

To obtain these fields, use either the command `class-param-list` or `class-param-get`

To filter the list, the CLI matches parameter values with those values specified on the command-line, only printing objects that match all of the specified constraints. For example:

```
1 xe vm-list HVM-boot-policy="BIOS order" power-state=halted
```

This command lists only those VMs for which *both* the field `power-state` has the value *halted* and the field `HVM-boot-policy` has the value *BIOS order*.

You can also filter the list by the value of keys in maps or by the existence of values in a set. The syntax for filtering based on keys in maps is `map-name:key=value`. The syntax for filtering based on values existing in a set is `set-name:contains=value`.

When scripting, a useful technique is passing `--minimal` on the command line, causing `xe` to print only the first field in a comma-separated list. For example, the command `xe vm-list --minimal` on a host with three VMs installed gives the three UUIDs of the VMs:

```
1      a85d6717-7264-d00e-069b-3b1d19d56ad9,aaa3eec5-9499-bcf3-4c03-  
      af10baea96b7, \  
2      42c044de-df69-4b30-89d9-2c199564581d
```

Secrets

XenServer provides a secrets mechanism to avoid passwords being stored in plaintext in command-line history or on API objects. XenCenter uses this feature automatically and it can also be used from the `xe` CLI for any command that requires a password.

Note:

Password secrets cannot be used to authenticate with a XenServer host from a remote instance of the `xe` CLI.

To create a secret object, run the following command on your XenServer host.

```
1 xe secret-create value=my-password
```

A secret is created and stored on the XenServer host. The command outputs the UUID of the secret object. For example, 99945d96-5890-de2a-3899-8c04ef2521db. Append `_secret` to the name of the password argument to pass this UUID to any command that requires a password.

Example: On the XenServer host where you created the secret, you can run the following command:

```
1  xe sr-create device-config:location=sr_address device-config:type=
   cifs device-config:username=cifs_username \
2  device-config:cifspassword_secret=secret_uuid name-label="CIFS ISO
   SR" type="iso" content-type="iso" shared="true"
```

Command history

Some `xe` commands, for example `xe vm-migrate` or `xe pool-enable-external-auth`, take secrets like passwords as parameters. These can end up in the shell history and during execution of the command are visible in the process table. It is therefore important to run these commands only in trustworthy environments.

For the bash shell, you can use the `HISTCONTROL` variable to control which commands are stored in the shell history.

xe command reference

This section groups the commands by the objects that the command addresses. These objects are listed alphabetically.

Appliance commands

Commands for creating and modifying VM appliances (also known as vApps). For more information, see [vApps](#).

Appliance parameters

Appliance commands have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The appliance uuid	Required
<code>name-description</code>	The appliance description	Optional

Parameter Name	Description	Type
<code>paused</code>		Optional
<code>force</code>	Force shutdown	Optional

appliance-assert-can-be-recovered

```
1 xe appliance-assert-can-be-recovered uuid=appliance-uuid database:vdi-uuid=vdi-uuid
```

Tests whether storage is available to recover this VM appliance/vApp.

appliance-create

```
1 xe appliance-create name=label=name=label [name-description=name-description]
```

Creates an appliance/vApp. For example:

```
1 xe appliance-create name=label=my_appliance
```

Add VMs to the appliance:

```
1 xe vm-param-set uuid=VM-UUID appliance=appliance-uuid
```

appliance-destroy

```
1 xe appliance-destroy uuid=appliance-uuid
```

Destroys an appliance/vApp. For example:

```
1 xe appliance-destroy uuid=appliance-uuid
```

appliance-recover

```
1 xe appliance-recover uuid=appliance-uuid database:vdi-uuid=vdi-uuid [
  paused=true | false]
```

Recover a VM appliance/vApp from the database contained in the supplied VDI.

appliance-shutdown

```
1 xe appliance-shutdown uuid=appliance-uuid [force=true | false]
```

Shuts down all VMs in an appliance/vApp. For example:

```
1 xe appliance-shutdown uuid=appliance-uuid
```

appliance-start

```
1 xe appliance-start uuid=appliance-uuid [paused=true | false]
```

Starts an appliance/vApp. For example:

```
1 xe appliance-start uuid=appliance-uuid
```

Audit commands

Audit commands download all of the available records of the RBAC audit file in the pool. If the optional parameter `since` is present, it downloads only the records from that specific point in time.

audit-log-get parameters

`audit-log-get` has the following parameters

Parameter Name	Description	Type
<code>filename</code>	Write the audit log of the pool to <i>file name</i>	Required
<code>since</code>	Specific date/time point	Optional

audit-log-get

```
1 xe audit-log-get [since=timestamp] filename=filename
```

For example, to obtain audit records of the pool since a precise millisecond timestamp, run the following command:

Run the following command:

```
1 xe audit-log-get since=2009-09-24T17:56:20.530Z filename=/tmp/auditlog-pool-actions.out
```

Bonding commands

Commands for working with network bonds, for resilience with physical interface failover. For more information, see [Networking](#).

The bond object is a reference object which glues together *master* and *member* PIFs. The master PIF is the bonding interface which must be used as the overall PIF to refer to the bond. The member PIFs are a set of two or more physical interfaces that have been combined into the high-level bonded interface.

Bond parameters

Bonds have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	Unique identifier/object reference for the bond	Read only
<code>master</code>	UUID for the main bond PIF	Read only
<code>members</code>	Set of UUIDs for the underlying bonded PIFs	Read only

`bond-create`

```
1 xe bond-create network-uuid=network_uuid pif-uuids=pif_uuid_1,  
  pif_uuid_2,...
```

Create a bonded network interface on the network specified from a list of existing PIF objects. The command fails in any of the following cases:

- If PIFs are in another bond already
- If any member has a VLAN tag set
- If the referenced PIFs are not on the same XenServer host
- If fewer than 2 PIFs are supplied

`bond-destroy`

```
1 xe bond-destroy uuid=bond_uuid
```

Deletes a bonded interface specified by its UUID from a host.

bond-set-mode

```
1 xe bond-set-mode uuid=bond_uuid mode=bond_mode
```

Change the bond mode.

CD commands

Commands for working with physical CD/DVD drives on XenServer hosts.

CD parameters

CDs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	Unique identifier/object reference for the CD	Read only
<code>name-label</code>	Name for the CD	Read/write
<code>name-description</code>	Description text for the CD	Read/write
<code>allowed-operations</code>	A list of the operations that can be performed on this CD	Read only set parameter
<code>current-operations</code>	A list of the operations that are currently in progress on this CD	Read only set parameter
<code>sr-uuid</code>	The unique identifier/object reference for the SR this CD is part of	Read only
<code>sr-name-label</code>	The name for the SR this CD is part of	Read only
<code>vbd-uuids</code>	A list of the unique identifiers for the VBDs on VMs that connect to this CD	Read only set parameter
<code>crashdump-uuids</code>	Not used on CDs. Because crashdumps cannot be written to CDs	Read only set parameter
<code>virtual-size</code>	Size of the CD as it appears to VMs (in bytes)	Read only

Parameter Name	Description	Type
<code>physical-utilisation</code>	Amount of physical space that the CD image takes up on the SR (in bytes)	Read only
<code>type</code>	Set to User for CDs	Read only
<code>sharable</code>	Whether or not the CD drive is sharable. Default is false .	Read only
<code>read-only</code>	Whether the CD is read-only, if false , the device is writable. Always true for CDs.	Read only
<code>storage-lock</code>	Value is true if this disk is locked at the storage level.	Read only
<code>parent</code>	Reference to the parent disk, if this CD is part of a chain.	Read only
<code>missing</code>	Value is true if SR scan operation reported this CD as not present on disk	Read only
<code>other-config</code>	A list of key-value pairs that specify extra configuration parameters for the CD	Read/write map parameter
<code>location</code>	The path on which the device is mounted	Read only
<code>managed</code>	Value is true if the device is managed	Read only
<code>xenstore-data</code>	Data to be inserted into the <code>xenstore</code> tree	Read only map parameter
<code>sm-config</code>	Names and descriptions of storage manager device config keys	Read only map parameter
<code>is-a-snapshot</code>	Value is true if this template is a CD snapshot	Read only
<code>snapshot_of</code>	The UUID of the CD that this template is a snapshot of	Read only
<code>snapshots</code>	The UUIDs of any snapshots that have been taken of this CD	Read only
<code>snapshot_time</code>	The timestamp of the snapshot operation	Read only

cd-list

```
1 xe cd-list [params=param1,param2,...] [parameter=parameter_value]
```

List the CDs and ISOs (CD image files) on the XenServer host or pool, filtering on the optional argument [params](#).

If the optional argument [params](#) is used, the value of [params](#) is a string containing a list of parameters of this object that you want to display. Alternatively, you can use the keyword [all](#) to show all parameters. When [params](#) is not used, the returned list shows a default subset of all available parameters.

Optional arguments can be any number of the [CD parameters](#) listed at the beginning of this section.

GFS2 cluster commands

Commands for working with GFS2 clustered pools.

GFS2 clustered pools are resource pools that have the clustering feature enabled. Use these pools with GFS2 SRs. For more information, see [GFS2 clustered pools](#)

The cluster and cluster-host objects can be listed with the standard object listing commands ([xe cluster-list](#) and [xe cluster-host-list](#)), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#).

Commands for working with GFS2 clustered pools.

GFS2 cluster parameters

GFS2 clusters have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the GFS2 cluster.	Read only
cluster-hosts	A list of unique identifiers/object references for the hosts in the GFS2 cluster.	Read only set parameter
cluster-token	The secret key used by xapi-clusterd when it talks to itself on other hosts.	Read only
cluster-stack	The technology stack providing the clustering capabilities. Possible values are corosync .	Read only

Parameter Name	Description	Type
<code>cluster-stack-version</code>	The version of the technology stack providing the clustering capabilities.	Read only
<code>allowed-operations</code>	Lists the operations allowed in this state. This list is advisory only and the GFS2 cluster state may have changed by the time a client reads this field.	Read only set parameter
<code>current-operations</code>	Lists the operations currently in process. This list is advisory only and the GFS2 cluster state may have changed by the time a client reads this field.	Read only set parameter
<code>token-timeout</code>	The <code>corosync</code> token timeout in seconds.	Read only
<code>token-timeout-coefficient</code>	The <code>corosync</code> token timeout coefficient in seconds.	Read only
<code>pool-auto-join</code>	Value is true if automatically joining new pool members to the GFS2 cluster. This is set to true .	Read only
<code>is-quorate</code>	Whether the GFS2 cluster has quorum.	Read only
<code>quorum</code>	The number of live hosts required for the GFS2 cluster to have quorum.	Read only
<code>live-hosts</code>	The number of hosts in the GFS2 cluster that are currently live.	Read only
<code>cluster-config</code>	A list of key/value pairs that specify extra configuration parameters for the GFS2 cluster.	Read only map parameter
<code>other-config</code>	A list of key/value pairs that specify extra configuration parameters for the GFS2 cluster.	Read/write map parameter

The hosts within the GFS2 cluster (`cluster-host`) have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the GFS2 cluster host. This is not the same as the host UUID.	Read only
<code>cluster</code>	The unique identifier/object reference for the GFS2 cluster this host is a member of.	Read only
<code>PIF</code>	The unique identifier/object reference for the PIF object.	Read only
<code>host</code>	The unique identifier/object reference for the host.	Read only
<code>enabled</code>	Whether the GFS2 cluster host believes that GFS2 clustering should be enabled on this host. This field can be altered by calling the enable/disable message on a GFS2 cluster host. Only enabled members run the underlying GFS2 cluster stack. Disabled members are still considered a member of the GFS2 cluster (see <code>joined</code>), and can be re-enabled by the user.	Read only
<code>joined</code>	Whether the GFS2 cluster host has joined the GFS2 cluster. Contrary to <code>enabled</code> , a host that is not joined is not considered a member of the GFS2 cluster, and hence <code>enable</code> and <code>disable</code> operations cannot be performed on this host.	Read only
<code>live</code>	Whether the cluster stack sees this host as live.	Read only

Parameter Name	Description	Type
<code>last-update-live</code>	Time when the live field was last updated based on information from the cluster stack.	Read only
<code>allowed-operations</code>	Lists the operations allowed in this state. This list is advisory only and the GFS2 cluster host state may have changed by the time a client reads this field.	Read only set parameter
<code>current-operations</code>	Lists the operations currently in process. This list is advisory only and the GFS2 cluster host state may have changed by the time a client reads this field.	Read only set parameter
<code>other-config</code>	A list of key/value pairs that specify extra configuration parameters for the GFS2 cluster host.	Read only map parameter

cluster-host-destroy

```
1 xe cluster-host-destroy uuid=host_uuid
```

Destroy a GFS2 cluster host, effectively leaving the GFS2 cluster.

cluster-host-disable

```
1 xe cluster-host-disable uuid=cluster_uuid
```

Disable GFS2 cluster membership for an enabled GFS2 cluster host.

cluster-host-enable

```
1 xe cluster-host-enable uuid=cluster_uuid
```

Enable GFS2 cluster membership for a disabled GFS2 cluster host.

cluster-host-force-destroy

```
1 xe cluster-host-force-destroy uuid=cluster_host
```

Destroy a GFS2 cluster host object forcefully, effectively leaving the GFS2 cluster.

cluster-pool-create

```
1 xe cluster-pool-create network-uuid=network_uuid [cluster-stack=
  cluster_stack] [token-timeout=token_timeout] [token-timeout-
  coefficient=token_timeout_coefficient]
```

Create pool-wide GFS2 cluster.

cluster-pool-destroy

```
1 xe cluster-pool-destroy cluster-uuid=cluster_uuid
```

Destroy pool-wide GFS2 cluster. The pool continues to exist, but it is no longer GFS2 clustered and can no longer use GFS2 SRs.

cluster-pool-force-destroy

```
1 xe cluster-pool-force-destroy cluster-uuid=cluster_uuid
```

Force destroy pool-wide GFS2 cluster.

cluster-pool-resync

```
1 xe cluster-pool-resync cluster-uuid=cluster_uuid
```

Resync a GFS2 cluster across a pool.

Console commands

Commands for working with consoles.

The console objects can be listed with the standard object listing command (`xe console-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#).

Console parameters

Consoles have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the console	Read only
<code>vm-uuid</code>	The unique identifier/object reference of the VM this console is open on	Read only
<code>vm-name-label</code>	The name of the VM this console is open on	Read only
<code>protocol</code>	Protocol this console uses. Possible values are <code>vt100</code> : VT100 terminal, <code>rfb</code> : Remote Framebuffer Protocol (as used in VNC), or <code>rdp</code> : Remote Desktop Protocol	Read only
<code>location</code>	URI for the console service	Read only
<code>other-config</code>	A list of key-value pairs that specify extra configuration parameters for the console.	Read/write map parameter

console

```
1 xe console
```

Attach to a particular console.

Diagnostic commands

Commands for gathering diagnostic information from XenServer.

diagnostic-compact

```
1 xe diagnostic-compact
```

Perform a major GC collection and heap compaction.

diagnostic-db-stats

```
1 xe diagnostic-db-stats
```

Print database statistics.

diagnostic-gc-stats

```
1 xe diagnostic-gc-stats
```

Print GC statistics.

diagnostic-license-status

```
1 xe diagnostic-license-status
```

Help diagnose pool-wide licensing problems.

diagnostic-net-stats

```
1 xe diagnostic-net-stats [uri=uri] [method=method] [params=param1,param2  
...]
```

Print network statistics.

diagnostic-timing-stats

```
1 xe diagnostic-timing-stats
```

Print timing statistics.

diagnostic-vdi-status

```
1 xe diagnostic-vdi-status uuid=vdi_uuid
```

Query the locking and sharing status of a VDI.

diagnostic-vm-status

```
1 xe diagnostic-vm-status uuid=vm_uuid
```

Query the hosts on which the VM can boot, check the sharing/locking status of all VBDs.

Disaster recovery commands

Commands for recovering VMs after a disaster

drtask-create

```
1 xe drtask-create type=type sr-whitelist=sr-white-list device-config=
   device-config
```

Creates a disaster recovery task. For example, to connect to an iSCSI SR in preparation for Disaster Recovery:

```
1 xe drtask-create type=lvmoiscsi device-config:target=target-ip-address
   \
2   device-config:targetIQN=targetIQN device-config:SCSIid=SCSIid \
3   sr-whitelist=sr-uuid-list
```

Note:

The command `sr-whitelist` lists SR UUIDs that are allowed. The `drtask-create` command only introduces and connects to an SR which has one of the allowed UUIDs

drtask-destroy

```
1 xe drtask-destroy uuid=dr-task-uuid
```

Destroys a disaster recovery task and forgets the introduced SR.

vm-assert-can-be-recovered

```
1 xe vm-assert-can-be-recovered uuid=vm-uuid database:vdi-uuid=vdi-uuid
```

Tests whether storage is available to recover this VM.

appliance-assert-can-be-recovered

```
1 xe appliance-assert-can-be-recovered uuid=appliance-uuid database:vdi-
   uuid=vdi-uuid
```

Checks whether the storage (containing the appliance's/vAPP disk) is visible.

appliance-recover

```
1 xe appliance-recover uuid=appliance-uuid database:vdi-uuid=vdi-uuid [
    force=true|false]
```

Recover an appliance/vAPP from the database contained in the supplied VDI.

vm-recover

```
1 xe vm-recover uuid=vm-uuid database:vdi-uuid=vdi-uuid [force=true|false
    ]
```

Recovers a VM from the database contained in the supplied VDI.

sr-enable-database-replication

```
1 xe sr-enable-database-replication uuid=sr_uuid
```

Enables XAPI database replication to the specified (shared) SR.

sr-disable-database-replication

```
1 xe sr-disable-database-replication uuid=sr_uuid
```

Disables XAPI database replication to the specified SR.

Example usage

The example below shows the DR CLI commands in context:

On the primary site, enable database replication:

```
1 xe sr-database-replication uuid=sr=uuid
```

After a disaster, on the secondary site, connect to the SR. The `device-config` command has the same fields as `sr-probe`.

```
1 xe drtask-create type=lvmoiscsi \
2     device-config:target=target ip address \
3     device-config:targetIQN=target-iqn \
4     device-config:SCSIid=scsi-id \
5     sr-whitelist=sr-uuid
```

Look for database VDIs on the SR:

```
1 xe vdi-list sr-uuid=sr-uuid type=Metadata
```

Query a database VDI for VMs present:

```
1 xe vm-list database:vdi-uuid=vdi-uuid
```

Recover a VM:

```
1 xe vm-recover uuid=vm-uuid database:vdi-uuid=vdi-uuid
```

Destroy the DR task. Any SRs introduced by the DR task and not required by VMs are destroyed:

```
1 xe drtask-destroy uuid=drtask-uuid
```

Event commands

Commands for working with events.

Event classes

Event classes are listed in the following table:

Class name	Description
<code>pool</code>	A pool of physical hosts
<code>vm</code>	A Virtual Machine
<code>host</code>	A physical host
<code>network</code>	A virtual network
<code>vif</code>	A virtual network interface
<code>pif</code>	A physical network interface (separate VLANs are represented as several PIFs)
<code>sr</code>	A storage repository
<code>vdi</code>	A virtual disk image
<code>vbd</code>	A virtual block device
<code>pbd</code>	The physical block devices through which hosts access SRs

`event-wait`

```
1 xe event-wait class=class_name [param-name=param_value] [param-name/=param_value]
```

Blocks other commands from running until an object exists that satisfies the conditions given on the command line. The argument `x=y` means “wait for field x to take value y” and `x/=y` means “wait for field x to take any value other than y.”

Example: wait for a specific VM to be running.

```
1 xe event-wait class=vm name=label=myvm power-state=running
```

Blocks other commands until a VM called `myvm` is in the `power-state` “running.”

Example: wait for a specific VM to reboot:

```
1 xe event-wait class=vm uuid=$VM start-time/= $(xe vm-list uuid=$VM  
  params=start-time --minimal)
```

Blocks other commands until a VM with UUID `$VM` reboots. The command uses the value of `start-time` to decide when the VM reboots.

The class name can be any of the [event classes](#) listed at the beginning of this section. The parameters can be any of the parameters listed in the CLI command `class-param-list`.

GPU commands

Commands for working with physical GPUs, GPU groups, and virtual GPUs.

The GPU objects can be listed with the standard object listing commands: `xe pgpu-list`, `xe gpu-group-list`, and `xe vgpu-list`. The parameters can be manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#).

Physical GPU parameters

Physical GPUS (pGPUs) have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the pGPU	Read only
<code>vendor-name</code>	The vendor name of the pGPU	Read only
<code>device-name</code>	The name assigned by the vendor to this pGPU model	Read only

Parameter Name	Description	Type
<code>gpu-group-uuid</code>	The unique identifier/object reference for the GPU group that this pGPU has been automatically assigned to by XenServer. Identical pGPUs across hosts in a pool are grouped	Read only
<code>gpu-group-name-label</code>	The name of the GPU group to which the pGPU is assigned	Read only
<code>host-uuid</code>	The unique identifier/object reference for the XenServer host to which the pGPU is connected	Read only
<code>host-name-label</code>	The name of the XenServer host to which the pGPU is connected	Read only
<code>pci-id</code>	PCI identifier	Read only
<code>dependencies</code>	Lists the dependent PCI devices passed-through to the same VM	Read/write map parameter
<code>other-config</code>	A list of key-value pairs that specify extra configuration parameters for the pGPU	Read/write map parameter
<code>supported-VGPU-types</code>	List of virtual GPU types supported by the underlying hardware	Read only
<code>enabled-VGPU-types</code>	List of virtual GPU types which have been enabled for this pGPU	Read/Write
<code>resident-VGPUs</code>	List of vGPUs running on this pGPU	Read only

pgpu-disable-dom0-access

```
1 xe pgpu-disable-dom0-access uuid=uuid
```

Disable PGPU access to dom0.

pgpu-enable-dom0-access

```
1 xe pgpu-enable-dom0-access uuid=uuid
```

Enable PGPU access to dom0.

GPU group parameters

GPU groups have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the GPU group	Read only
<code>name-label</code>	The name of the GPU group	Read/write
<code>name-description</code>	The descriptive text of the GPU group	Read/write
<code>VGPU-uuids</code>	Lists the unique identifier/object references for the virtual GPUs in the GPU group	Read only set parameter
<code>PGPU-uuids</code>	Lists the unique identifier/object references for the pGPUs in the GPU group	Read only set parameter
<code>other-config</code>	A list of key-value pairs that specify extra configuration parameters for the GPU group	Read/write map parameter
<code>supported-VGPU-types</code>	Union of all virtual GPU types supported by the underlying hardware	Read only
<code>enabled-VGPU-types</code>	Union of all virtual GPU types which have been enabled on the underlying pGPUs	Read only
<code>allocation-algorithm</code>	Depth-first/Breadth-first setting for allocation virtual GPUs on pGPUs within the group	Read/write enum parameter

GPU group operations Commands for working with GPU Groups

gpu-group-create

```
1 xe gpu-group-create name=label=name_for_group [name-description=
    description]
```

Creates a new (empty) GPU Group into which pGPUs can be moved.

gpu-group-destroy

```
1 xe gpu-group-destroy uuid=uuid_of_group
```

Destroys the GPU Group; only permitted for empty groups.

gpu-group-get-remaining-capacity

```
1 xe gpu-group-get-remaining-capacity uuid=uuid_of_group vgpu-type-uuid=
    uuid_of_vgpu_type
```

Returns how many more virtual GPUs of the specified type can be instantiated in this GPU Group.

gpu-group-param-set

```
1 xe gpu-group-param-set uuid=uuid_of_group allocation-algorithm=breadth-
    first|depth-first
```

Changes the algorithm that the GPU group uses to allocate virtual GPUs to pGPUs.

Virtual GPU parameters

Virtual GPUs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the virtual GPU	Read only
<code>vm-uuid</code>	The unique identifier/object reference for the VM to which the virtual GPU is assigned	Read only
<code>vm-name=label</code>	The name of the VM to which the virtual GPU is assigned	Read only
<code>gpu-group-uuid</code>	The unique identifier/object reference for the GPU group in which the virtual GPU is contained	Read only
<code>gpu-group-name=label</code>	The name of the GPU group in which the virtual GPU is contained	Read only

Parameter Name	Description	Type
<code>currently-attached</code>	Value is true if a VM with GPU pass-through is running, false otherwise	Read only
<code>other-config</code>	A list of key-value pairs that specify extra configuration parameters for the virtual GPU	Read/write map parameter
<code>type-uuid</code>	The unique identifier/object reference for the virtual GPU type of this virtual GPU	Read/write map parameter
<code>type-model-name</code>	Model name associated with the virtual GPU type	Read only

Virtual GPU type parameters

Note:

GPU pass-through and virtual GPUs are not compatible with live migration, storage live migration, or VM Suspend unless supported software and graphics cards from GPU vendors are present. VMs without this support cannot be migrated to avoid downtime. For information about NVIDIA vGPU compatibility with live migration, storage live migration, and VM Suspend, see [Graphics](#).

Virtual GPU Types have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the virtual GPU type	Read only
<code>vendor-name</code>	Name of virtual GPU vendor	Read only
<code>model-name</code>	Model name associated with the virtual GPU type	Read only
<code>freeze-frame</code>	Frame buffer size of the virtual GPU type, in bytes	Read only
<code>max-heads</code>	Maximum number of displays supported by the virtual GPU type	Read only
<code>supported-on-PGPUs</code>	List of pGPUs that support this virtual GPU type	Read only

Parameter Name	Description	Type
<code>enabled-on-PGPUs</code>	List of pGPUs that have this virtual GPU type enabled	Read only
<code>VGPU-uuids</code>	List of virtual GPUs of this type	Read only

Virtual GPU operations

`vgpu-create`

```
1 xe vgpu-create vm-uuid=uuid_of_vm gpu_group_uuid=uuid_of_gpu_group [
  vgpu-type-uuid=uuid_of_vgpu-type]
```

Creates a virtual GPU. This command attaches the VM to the specified GPU group and optionally specifies the virtual GPU type. If no virtual GPU type is specified, the ‘pass-through’ type is assumed.

`vgpu-destroy`

```
1 xe vgpu-destroy uuid=uuid_of_vgpu
```

Destroy the specified virtual GPU.

Disabling VNC for VMs with virtual GPU

```
1 xe vm-param-add uuid=uuid_of_vmparam-name=platform vgpu_vnc_enabled=
  true|false
```

Using **false** disables the VNC console for a VM as it passes `disablevnc=1` through to the display emulator. By default, VNC is enabled.

Host commands

Commands for interacting with XenServer host.

XenServer hosts are the physical servers running XenServer software. They have VMs running on them under the control of a special privileged Virtual Machine, known as the control domain or domain 0.

The XenServer host objects can be listed with the standard object listing commands: `xe host-list`, `xe host-cpu-list`, and `xe host-crashdump-list`). The parameters can be manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#).

Host selectors

Several of the commands listed here have a common mechanism for selecting one or more XenServer hosts on which to perform the operation. The simplest is by supplying the argument

`host=uuid_or_name_label`. You can also specify XenServer by filtering the full list of hosts on the values of fields. For example, specifying `enabled=true` selects all XenServer hosts whose `enabled` field is equal to `true`. Where multiple XenServer hosts match and the operation can be performed on multiple XenServer hosts, you must specify `--multiple` to perform the operation. The full list of parameters that can be matched is described at the beginning of this section. You can obtain this list of commands by running the command `xe host-list params=all`. If no parameters to select XenServer hosts are given, the operation is performed on all XenServer hosts.

Host parameters

XenServer hosts have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the XenServer host	Read only
<code>name-label</code>	The name of the XenServer host	Read/write
<code>name-description</code>	The description string of the XenServer host	Read only
<code>enabled</code>	Value is <code>false</code> if disabled. This prevents any new VMs from starting on the hosts and prepares the hosts to be shut down or rebooted. Value is <code>true</code> if the host is enabled	Read only
<code>API-version-major</code>	Major version number	Read only
<code>API-version-minor</code>	Minor version number	Read only
<code>API-version-vendor</code>	Identification of API vendor	Read only
<code>API-version-vendor-implementation</code>	Details of vendor implementation	Read only map parameter
<code>logging</code>	Logging configuration	Read/write map parameter
<code>suspend-image-sr-uuid</code>	The unique identifier/object reference for the SR where suspended images are put	Read/write
<code>crash-dump-sr-uuid</code>	The unique identifier/object reference for the SR where crash dumps are put	Read/write

Parameter Name	Description	Type
<code>software-version</code>	List of versioning parameters and their values	Read only map parameter
<code>capabilities</code>	List of Xen versions that the XenServer host can run	Read only set parameter
<code>ntp-mode</code>	Selects NTP mode • Factory • DHCP • Custom • Disabled	Read/write
<code>ntp-custom-servers</code>	Comma-separated custom NTP server list	Read/write (set)
<code>timezone</code>	IANA timezone identifier	Read/write
<code>max-cstate</code>	Maximum CPU C-state: “” (unlimited), 0 (C0), 1 (C1)	Read/write
<code>other-config</code>	A list of key-value pairs that specify extra configuration parameters for the XenServer host	Read/write map parameter
<code>chipset-info</code>	A list of key-value pairs that specify information about the chipset	Read only map parameter
<code>hostname</code>	XenServer host host name	Read only
<code>address</code>	XenServer host IP address	Read only
<code>license-server</code>	A list of key-value pairs that specify information about the license server. The default port for communications with the Citrix License Server and License Activation Service is 8083. For information on changing port numbers due to conflicts, see Change port numbers	Read only map parameter

Parameter Name	Description	Type
<code>supported-bootloaders</code>	List of bootloaders that the XenServer host supports, for example, <code>pygrub</code> , <code>eliloader</code>	Read only set parameter
<code>memory-total</code>	Total amount of physical RAM on the XenServer host, in bytes	Read only
<code>memory-free</code>	Total amount of physical RAM remaining that can be allocated to VMs, in bytes	Read only
<code>host-metrics-live</code>	Value is true if the host is operational, false otherwise	Read only
<code>logging</code>	The <code>syslog_destination</code> key can be set to the host name of a remote listening syslog service.	Read/write map parameter
<code>allowed-operations</code>	Lists the operations allowed in this state. This list is advisory only and the host state may have changed by the time a client reads this field.	Read only set parameter
<code>current-operations</code>	Lists the operations currently in process. This list is advisory only and the host state may have changed by the time a client reads this field.	Read only set parameter
<code>patches</code>	Set of host patches	Read only set parameter
<code>blobs</code>	Binary data store	Read only
<code>memory-free-computed</code>	A conservative estimate of the maximum amount of memory free on a host	Read only
<code>ha-statefiles</code>	The UUIDs of all HA state files	Read only
<code>ha-network-peers</code>	The UUIDs of all hosts that can host the VMs on this host if there is a failure	Read only
<code>external-auth-type</code>	Type of external authentication, for example, Active Directory.	Read only

Parameter Name	Description	Type
<code>external-auth-service-name</code>	The name of the external authentication service	Read only
<code>external-auth-configuration</code>	Configuration information for the external authentication service.	Read only map parameter

Examples:

- `xe host-param-set uuid=<uuid> ntp-mode=Factory`
- `xe host-param-set uuid=<uuid> ntp-custom-servers=server1,server2`
- `xe host-param-set uuid=<uuid> timezone=UTC`
- `xe host-param-set uuid=<uuid> max-cstate=1`

XenServer hosts contain some other objects that also have parameter lists.

CPUs on XenServer hosts have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the CPU	Read only
<code>number</code>	The number of the physical CPU core within the XenServer host	Read only
<code>vendor</code>	The vendor string for the CPU name	Read only
<code>speed</code>	The CPU clock speed, in Hz	Read only
<code>modelname</code>	The vendor string for the CPU model, for example, “Intel(R) Xeon(TM) CPU 3.00 GHz”	Read only
<code>stepping</code>	The CPU revision number	Read only
<code>flags</code>	The flags of the physical CPU (a decoded version of the features field)	Read only
<code>Utilisation</code>	The current CPU utilization	Read only
<code>host-uuid</code>	The UUID of the host the CPU is in	Read only

Parameter Name	Description	Type
<code>model</code>	The model number of the physical CPU	Read only
<code>family</code>	The physical CPU family number	Read only

Crash dumps on XenServer hosts have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the crashdump	Read only
<code>host</code>	XenServer host the crashdump corresponds to	Read only
<code>timestamp</code>	Timestamp of the date and time that the crashdump occurred, in the form <code>yyyymmdd-hhmmss-ABC</code> , where <i>ABC</i> is the timezone indicator, for example, GMT	Read only
<code>size</code>	Size of the crashdump, in bytes	Read only

host-all-editions

```
1 xe host-all-editions
```

Get a list of all available editions

host-apply-edition

```
1 xe host-apply-edition [host-uuid=host_uuid] [edition=premium-per-socket
| standard-per-socket | cloud]
```

Assigns the XenServer license to a host server. When you assign a license, XenServer contacts the license server and requests the specified type of license. If a license is available, it is then checked out from the license server.

For initial licensing configuration, see also `license-server-address` and `license-server-port`.

Edition values include “cloud”, “free”, “premium-per-socket”, and “standard-per-socket”. In future, only “cloud” is supported.

host-backup

```
1 xe host-backup file-name=backup_filename host=host_name
```

Download a backup of the control domain of the specified XenServer host to the machine that the command is invoked from. Save it there as a file with the name `file-name`.

Important:

While the `xe host-backup` command works if run on the local host (that is, without a specific host name specified), do not use it this way. Doing so would fill up the control domain partition with the backup file. Only use the command from a remote off-host machine where you have space to hold the backup file.

host-bugreport-upload

```
1 xe host-bugreport-upload [host-selector=host_selector_value...] [url=
  destination_url http-proxy=http_proxy_name]
```

Generate a fresh bug report (using `xen-bugtool`, with all optional files included) and upload to the Support FTP site or some other location.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

Optional parameters are `http-proxy`: use specified HTTP proxy, and `url`: upload to this destination URL. If optional parameters are not used, no proxy server is identified and the destination is the default Support FTP site.

host-call-plugin

```
1 xe host-call-plugin host-uuid=host_uuid plugin=plugin fn=function [args
  =args]
```

Calls the function within the plug-in on the given host with optional arguments.

host-compute-free-memory

```
1 xe host-compute-free-memory
```

Computes the amount of free memory on the host.

host-compute-memory-overhead

```
1 xe host-compute-memory-overhead
```

Computes the virtualization memory overhead of a host.

host-cpu-info

```
1 xe host-cpu-info [uuid=uuid]
```

Lists information about the host's physical CPUs.

host-crashdump-destroy

```
1 xe host-crashdump-destroy uuid=crashdump_uuid
```

Delete a host crashdump specified by its UUID from the XenServer host.

host-crashdump-upload

```
1 xe host-crashdump-upload uuid=crashdump_uuid [url=destination_url] [
    http-proxy=http_proxy_name]
```

Upload a crashdump to the Support FTP site or other location. If optional parameters are not used, no proxy server is identified and the destination is the default Support FTP site. Optional parameters are [http-proxy](#): use specified HTTP proxy, and [url](#): upload to this destination URL.

host-declare-dead

```
1 xe host-declare-dead uuid=host_uuid
```

Declare that the host is dead without contacting it explicitly.

Warning:

This call is dangerous and can cause data loss if the host is not actually dead.

host-disable

```
1 xe host-disable [host-selector=host_selector_value...]
```

Disables the specified XenServer hosts, which prevents any new VMs from starting on them. This action prepares the XenServer hosts to be shut down or rebooted. After that host reboots, if all conditions for enabling are met (for example, storage is available), the host is automatically re-enabled.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

host-disable-display

```
1 xe host-disable-display uuid=host_uuid
```

Disable display for the host.

host-disable-local-storage-caching

```
1 xe host-disable-local-storage-caching
```

Disable local storage caching on the specified host.

host-dmesg

```
1 xe host-dmesg [host-selector=host_selector_value...]
```

Get a Xen [dmesg](#) (the output of the kernel ring buffer) from specified XenServer hosts.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

host-emergency-ha-disable

```
1 xe host-emergency-ha-disable [--force]
```

Disable HA on the local host. Only to be used to recover a pool with a broken HA setup.

host-emergency-management-reconfigure

```
1 xe host-emergency-management-reconfigure interface=  
    uuid_of_management_interface_pif
```

Reconfigure the management interface of this XenServer host. Use this command only if the XenServer host is in emergency mode. Emergency mode means that the host is a member in a resource pool whose pool coordinator has disappeared from the network and cannot be contacted after a number of retries.

host-emergency-reset-server-certificate

```
1 xe host-emergency-reset-server-certificate
```

Installs a self-signed certificate on the XenServer host where the command is run.

host-enable

```
1 xe host-enable [host-selector=host_selector_value...]
```

Enables the specified XenServer hosts, which allows new VMs to be started on them.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

host-enable-display

```
1 xe host-enable-display uuid=host_uuid
```

Enable display for the host.

host-enable-local-storage-caching

```
1 xe host-enable-local-storage-caching sr-uuid=sr_uuid
```

Enable local storage caching on the specified host.

host-evacuate

```
1 xe host-evacuate [host-selector=host_selector_value...]
```

Live migrates all running VMs to other suitable hosts on a pool. First, disable the host by using the `host-disable` command.

If the evacuated host is the pool coordinator, then another host must be selected to be the pool coordinator. To change the pool coordinator with HA disabled, use the `pool-designate-new-master` command. For more information, see [pool-designate-new-master](#).

With HA enabled, your only option is to shut down the XenServer host, which causes HA to elect a new pool coordinator at random. For more information, see [host-shutdown](#).

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

host-forget

```
1 xe host-forget uuid=host_uuid
```

The XAPI agent forgets about the specified XenServer host without contacting it explicitly.

Use the `--force` parameter to avoid being prompted to confirm that you really want to perform this operation.

Warning:

Do not use this command if HA is enabled on the pool. Disable HA first, then enable it again after you've forgotten the host.

This command is useful if the XenServer host to “forget” is dead. However, if the XenServer host is live and part of the pool, use `xe pool-eject` instead.

host-get-server-certificate

```
1 xe host-get-server-certificate
```

Get the installed server TLS certificate.

host-get-sm-diagnostics

```
1 xe host-get-sm-diagnostics uuid=uuid
```

Display per-host SM diagnostic information.

host-get-system-status

```
1 xe host-get-system-status filename=name_for_status_file [entries=  
    comma_separated_list] [output=tar.bz2|zip] [host-selector=  
    host_selector_value...]
```

Download system status information into the specified file. The optional parameter `entries` is a comma-separated list of system status entries, taken from the capabilities XML fragment returned by the `host-get-system-status-capabilities` command. For more information, see [host-get-system-status-capabilities](#). If not specified, all system status information is saved in the file. The parameter `output` may be `tar.bz2` (the default) or `zip`. If this parameter is not specified, the file is saved in `tar.bz2` form.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above).

host-get-system-status-capabilities

```
1 xe host-get-system-status-capabilities [host-selector=  
    host_selector_value...]
```

Get system status capabilities for the specified hosts. The capabilities are returned as an XML fragment that similar to the following example:

```
1 <?xml version="1.0" ?>  
2 <system-status-capabilities>  
3   <capability content-type="text/plain" default-checked="yes" key="  
4     xenserver-logs" \  
5     max-size="150425200" max-time="-1" min-size="150425200" min-  
6     time="-1" \  
7     pii="maybe"/>  
8   <capability content-type="text/plain" default-checked="yes" \  
9     key="xenserver-install" max-size="51200" max-time="-1" min-size  
10    ="10240" \  
    min-time="-1" pii="maybe"/>  
    ...  
</system-status-capabilities>
```

Each capability entity can have the following attributes.

- **key** A unique identifier for the capability.
- **content-type** Can be either `text/plain` or `application/data`. Indicates whether a UI can render the entries for human consumption.
- **default-checked** Can be either `yes` or `no`. Indicates whether a UI selects this entry by default.

- `min-size`, `max-size` Indicates an approximate range for the size, in bytes, of this entry. -1 indicates that the size is unimportant.
- `min-time`, `max-time` Indicate an approximate range for the time, in seconds, taken to collect this entry. -1 indicates that the time is unimportant.
- `pii` Personally identifiable information. Indicates whether the entry has information that can identify the system owner or details of their network topology. The attribute can have one of the following values:
 - `no`: no PII is in these entries
 - `yes`: PII likely or certainly is in these entries
 - `maybe`: you might want to audit these entries for PII
 - `if_customized` if the files are unmodified, then they contain no PII. However, because we encourage editing of these files, PII might have been introduced by such customization. This value is used in particular for the networking scripts in the control domain.

Passwords are never to be included in any bug report, regardless of any PII declaration.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above).

host-get-thread-diagnostics

```
1 xe host-get-thread-diagnostics uuid=uuid
```

Display per-host thread diagnostic information.

host-get-vms-which-prevent-evacuation

```
1 xe host-get-vms-which-prevent-evacuation uuid=uuid
```

Return a list of VMs which prevent the evacuation of a specific host and display reasons for each one.

host-is-in-emergency-mode

```
1 xe host-is-in-emergency-mode
```

Returns **true** if the host the CLI is talking to is in emergency mode, **false** otherwise. This CLI command works directly on pool member hosts even with no pool coordinator present.

host-license-view

```
1 xe host-license-view [host-uuid=host_uuid]
```

Displays the contents of the XenServer host license.

host-logs-download

```
1 xe host-logs-download [file-name=logfile_name] [host-selector=  
  host_selector_value...]
```

Download a copy of the logs of the specified XenServer hosts. The copy is saved by default in a time-stamped file named `hostname-yyyy-mm-dd T hh:mm:ssZ.tar.gz`. You can specify a different file name using the optional parameter *file-name*.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

Important:

While the `xe host-logs-download` command works if run on the local host (that is, without a specific host name specified), do *not* use it this way. Doing so clutters the control domain partition with the copy of the logs. Only use the command from a remote off-host machine where you have space to hold the copy of the logs.

host-management-disable

```
1 xe host-management-disable
```

Disables the host agent listening on an external management network interface and disconnects all connected API clients (such as the XenCenter). This command operates directly on the XenServer host the CLI is connected to. The command is not forwarded to the pool coordinator when applied to a member XenServer host.

Warning:

Be careful when using this CLI command off-host. After this command is run, you cannot connect to the control domain remotely over the network to re-enable the host agent.

host-management-reconfigure

```
1 xe host-management-reconfigure [interface=device] [pif-uuid=uuid]
```

Reconfigures the XenServer host to use the specified network interface as its management interface, which is the interface that is used to connect to the XenCenter. The command rewrites the `MANAGEMENT_INTERFACE` key in `/etc/xensource-inventory`.

If the device name of an interface (which must have an IP address) is specified, the XenServer host immediately rebinds. This command works both in normal and emergency mode.

If the UUID of a PIF object is specified, the XenServer host determines which IP address to rebind to itself. It must not be in emergency mode when this command is run.

Warning:

Be careful when using this CLI command off-host and ensure that you have network connectivity on the new interface. Use `xe pif-reconfigure` to set one up first. Otherwise, subsequent CLI commands are unable to reach the XenServer host.

host-power-on

```
1 xe host-power-on [host=host_uuid]
```

Turns on power on XenServer hosts with the *Host Power On* function enabled. Before using this command, enable `host-set-power-on` on the host.

host-reboot

```
1 xe host-reboot [host-selector=host_selector_value...]
```

Reboot the specified XenServer hosts. The specified hosts must be disabled first using the `xe host-disable` command, otherwise a `HOST_IN_USE` error message is displayed.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

If the specified XenServer hosts are members of a pool, the loss of connectivity on shutdown is handled and the pool recovers when the XenServer hosts return. The other members and the pool coordinator continue to function.

If you shut down the pool coordinator, the pool is out of action until one of the following actions occurs:

- You make one of the members into the pool coordinator
- The original pool coordinator is rebooted and back on line.

When the pool coordinator is back online, the members reconnect and synchronize with the pool coordinator.

host-restore

```
1 xe host-restore [file-name=backup_filename] [host-selector=  
    host_selector_value...]
```

Restore a backup named `file-name` of the XenServer host control software. The use of the word “restore” here does not mean a full restore in the usual sense, it merely means that the compressed backup file has been uncompressed and unpacked onto the secondary partition. After you’ve done a `xe host-restore`, you have to boot the Install CD and use its Restore from Backup option.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

host-send-debug-keys

```
1 xe host-send-debug-keys host-uuid=host_uuid keys=keys
```

Send specified hypervisor debug keys to specified host.

host-server-certificate-install

```
1 xe host-server-certificate-install certificate=path_to_certificate_file  
    private-key=path_to_private_key [certificate-chain=  
    path_to_chain_file] [host=host_name | uuid=host_uuid]
```

Install a TLS certificate on a XenServer host.

host-set-hostname-live

```
1 xe host-set-hostname-live host-uuid=uuid_of_host host-name=new_hostname
```

Change the host name of the XenServer host specified by `host-uuid`. This command persistently sets both the host name in the control domain database and the actual Linux host name of the XenServer host. The value of `host-name` is *not* the same as the value of the `name_label` field.

host-set-power-on-mode

```
1 xe host-set-power-on-mode host=host_uuid power-on-mode={  
2   "" | "wake-on-lan" | "IPMI" | "custom" }  
3   \
```

```
4 [ power-on-config:power_on_ip=ip-address power-on-config:  
    power_on_user=user power-on-config:power_on_password_secret=  
    secret-uuid ]
```

Use to enable the *Host Power On* function on XenServer hosts that are compatible with remote power solutions. When using the `host-set-power-on` command, you must specify the type of power management solution on the host (that is, the power-on-mode). Then specify configuration options using the power-on-config argument and its associated key-value pairs.

To use the secrets feature to store your password, specify the key "`power_on_password_secret`". For more information, see [Secrets](#).

host-shutdown

```
1 xe host-shutdown [host-selector=host_selector_value...]
```

Shut down the specified XenServer hosts. The specified XenServer hosts must be disabled first using the `xe host-disable` command, otherwise a `HOST_IN_USE` error message is displayed.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

If the specified XenServer hosts are members of a pool, the loss of connectivity on shutdown is handled and the pool recovers when the XenServer hosts returns. The other members and the pool coordinator continue to function.

If you shut down the pool coordinator, the pool is out of action until one of the following actions occurs:

- You make one of the members into the pool coordinator
- The original pool coordinator is rebooted and back on line.

When the pool coordinator is back online, the members reconnect and synchronize with the pool coordinator.

If HA is enabled for the pool, one of the members is made into a pool coordinator automatically. If HA is disabled, you must manually designate the desired XenServer host as pool coordinator with the `pool-designate-new-master` command. For more information, see [pool-designate-new-master](#).

host-sm-dp-destroy

```
1 xe host-sm-dp-destroy uuid=uuid dp=dp [allow-leak=true|false]
```

Attempt to destroy and clean up a storage datapath on a host. If `allow-leak=true` is provided then it deletes all records of the datapath even if it is not shut down cleanly.

host-sync-data

```
1 xe host-sync-data
```

Synchronize the data stored on the pool coordinator with the named host. This does not include the database data).

host-syslog-reconfigure

```
1 xe host-syslog-reconfigure [host-selector=host_selector_value...]
```

Reconfigure the `syslog` daemon on the specified XenServer hosts. This command applies the configuration information defined in the host `logging` parameter.

The hosts on which to perform this operation are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

host-data-source-list

```
1 xe host-data-source-list [host-selectors=host selector value...]
```

List the data sources that can be recorded for a host.

Select the hosts on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all hosts.

Data sources have two parameters –`standard` and `enabled`. This command outputs the values of the parameters:

- If a data source has `enabled` set to **true**, the metrics are currently being recorded to the performance database.
- If a data source has `standard` set to **true**, the metrics are recorded to the performance database *by default*. The value of `enabled` is also set to **true** for this data source.
- If a data source has `standard` set to **false**, the metrics are *not* recorded to the performance database by default. The value of `enabled` is also set to **false** for this data source.

To start recording data source metrics to the performance database, run the `host-data-source-record` command. This command sets `enabled` to **true**. To stop, run the `host-data-source-forget`. This command sets `enabled` to **false**.

host-data-source-record

```
1 xe host-data-source-record data-source=name_description_of_data_source  
   [host-selectors=host_selector_value...]
```

Record the specified data source for a host.

This operation writes the information from the data source to the persistent performance metrics database of the specified hosts. For performance reasons, this database is distinct from the normal agent database.

Select the hosts on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all hosts.

host-data-source-forget

```
1 xe host-data-source-forget data-source=name_description_of_data_source  
   [host-selectors=host_selector_value...]
```

Stop recording the specified data source for a host and forget all of the recorded data.

Select the hosts on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all hosts.

host-data-source-query

```
1 xe host-data-source-query data-source=name_description_of_data_source [  
   host-selectors=host_selector_value...]
```

Display the specified data source for a host.

Select the hosts on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all hosts.

Message commands

Commands for working with messages. Messages are created to notify users of significant events, and are displayed in XenCenter as alerts.

The message objects can be listed with the standard object listing command (`xe message-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

Message parameters

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the message	Read only
<code>name</code>	The unique name of the message	Read only
<code>priority</code>	The message priority. Higher numbers indicate greater priority	Read only
<code>class</code>	The message class, for example VM.	Read only
<code>obj-uuid</code>	The uuid of the affected object.	Read only
<code>timestamp</code>	The time that the message was generated.	Read only
<code>body</code>	The message content.	Read only

message-create

```
1 xe message-create name=message_name body=message_text [[host-uuid=
    uuid_of_host] | [sr-uuid=uuid_of_sr] | [vm-uuid=uuid_of_vm] | [pool-
    uuid=uuid_of_pool]]
```

Creates a message.

message-destroy

```
1 xe message-destroy [uuid=message_uuid]
```

Destroys an existing message. You can build a script to destroy all messages. For example:


```
1 # Dismiss all alerts \
2   IFS=","; for m in $(xe message-list params=uuid --minimal); do \
3     xe message-destroy uuid=$m \
4     done
```

Network commands

Commands for working with networks.

The network objects can be listed with the standard object listing command (`xe network-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

Network parameters

Networks have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the network	Read only
<code>name-label</code>	The name of the network	Read/write
<code>name-description</code>	The description text of the network	Read/write
<code>VIF-uuids</code>	A list of unique identifiers of the VIFs (virtual network interfaces) that are attached from VMs to this network	Read only set parameter
<code>PIF-uuids</code>	A list of unique identifiers of the PIFs (physical network interfaces) that are attached from XenServer hosts to this network	Read only set parameter
<code>bridge</code>	Name of the bridge corresponding to this network on the local XenServer host	Read only

Parameter Name	Description	Type
<code>default-locking-mode</code>	A network object used with VIF objects for ARP filtering. Set to <code>unlocked</code> to remove all the filtering rules associated with the VIF. Set to <code>disabled</code> so the VIF drops all traffic.	Read/write
<code>purpose</code>	Set of purposes for which the XenServer host uses this network. Set to <code>nbd</code> to use the network to make NBD connections.	Read/write
<code>other-config:staticroutes</code>	Comma-separated list of <i>subnet/netmask/gateway</i> formatted entries specifying the gateway address through which to route subnets. For example, setting <code>other-config:static-routes</code> to <code>172.16.0.0/15/192.168.0.3,172.18.0.0/16/192.168.0.4</code> causes traffic on <code>172.16.0.0/15</code> to be routed over <code>192.168.0.3</code> and traffic on <code>172.18.0.0/16</code> to be routed over <code>192.168.0.4</code> .	Read/write
<code>other-config:ethtoolautoneg</code>	Set to <code>no</code> to disable autonegotiation of the physical interface or bridge. Default is <code>yes</code> .	Read/write
<code>other-config:ethtool-rx</code>	Set to <code>on</code> to enable receive checksum, off to disable	Read/write
<code>other-config:ethtool-tx</code>	Set to <code>on</code> to enable transmit checksum, off to disable	Read/write
<code>other-config:ethtool-sg</code>	Set to <code>on</code> to enable scatter gather, off to disable	Read/write
<code>other-config:ethtool-tso</code>	Set to <code>on</code> to enable TCP segmentation offload, off to disable	Read/write

Parameter Name	Description	Type
<code>other-config:ethtool-ufo</code>	Set to on to enable UDP fragment offload, off to disable	Read/write
<code>other-config:ethtool-gso</code>	Set to on to enable generic segmentation offload, off to disable	Read/write
<code>blobs</code>	Binary data store	Read only

network-create

```
1 xe network-create name=label=name_for_network [name-description=
  descriptive_text]
```

Creates a network.

network-destroy

```
1 xe network-destroy uuid=network_uuid
```

Destroys an existing network.

SR-IOV commands

Commands for working with SR-IOV.

The `network-sriov` objects can be listed with the standard object listing command (`xe network-sriov-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

SR-IOV parameters

SR-IOV has the following parameters:

Parameter Name	Description	Type
<code>physical-PIF</code>	The PIF to enable SR-IOV.	Read only

Parameter Name	Description	Type
<code>logical-pif</code>	An SR-IOV logical PIF. Users can use this parameter to create an SR-IOV VLAN network.	Read only
<code>requires-reboot</code>	If set to true, used to reboot host to bring SR-IOV enabling into effect.	Read only
<code>remaining-capacity</code>	Number of available VFs remaining.	Read only

network-sriov-create

```
1 xe network-sriov-create network-uuid=network_uuid pif-uuid=
   physical_pif_uuid
```

Creates an SR-IOV network object for a given physical PIF and enables SR-IOV on the physical PIF.

network-sriov-destroy

```
1 xe network-sriov-destroy uuid=network_sriov_uuid
```

Removes a network SR-IOV object and disables SR-IOV on its physical PIF.

Assign an SR-IOV VF

```
1 xe vif-create device=device_index mac=vf_mac_address network-uuid=
   sriov_network vm-uuid=vm_uuid
```

Assigns a VF from an SR-IOV network to a VM.

SDN Controller commands

Commands for working with the SDN controller.

sdn-controller-forget

```
1 xe sdn-controller-introduce [address=address] [protocol=protocol] [tcp-
   port=tcp_port]
```

Introduce an SDN controller.

sdn-controller-introduce

```
1 xe sdn-controller-forget uuid=uuid
```

Remove an SDN controller.

Tunnel commands

Commands for working with tunnels.

tunnel-create

```
1 xe tunnel-create pif-uuid=pif_uuid network-uuid=network_uuid
```

Create a new tunnel on a host.

tunnel-destroy

```
1 xe tunnel-destroy uuid=uuid
```

Destroy a tunnel.

Patch commands

Commands for working with patches.

patch-apply

```
1 xe patch-apply uuid=patch_uuid host-uuid=host_uuid
```

Apply the previously uploaded patch to the specified host.

patch-clean

```
1 xe patch-clean uuid=uuid
```

Delete a previously uploaded patch file.

patch-destroy

```
1 xe patch-destroy uuid=uuid
```

Remove an unapplied patch record and files from the host.

patch-pool-apply

```
1 xe patch-pool-apply uuid=uuid
```

Apply the previously uploaded patch to all hosts in the pool.

patch-pool-clean

```
1 xe patch-pool-clean uuid=uuid
```

Delete a previously uploaded patch file on all hosts in the pool.

patch-precheck

```
1 xe patch-precheck uuid=uuid host-uuid=host_uuid
```

Run the pre-checks contained within the patch previously uploaded to the specified host.

patch-upload

```
1 xe patch-upload file-name=file_name
```

Upload a patch file to the host.

PBD commands

Commands for working with PBDs (Physical Block Devices). PBDs are the software objects through which the XenServer host accesses storage repositories (SRs).

The PBD objects can be listed with the standard object listing command (`xe pbd-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

PBD parameters

PBDs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the PBD.	Read only
<code>sr-uuid</code>	The storage repository that the PBD points to	Read only
<code>device-config</code>	Extra configuration information that is provided to the SR-backend-driver of a host	Read only map parameter
<code>currently-attached</code>	Value is true if the SR is attached on this host, false otherwise	Read only
<code>host-uuid</code>	UUID of the physical machine on which the PBD is available	Read only
<code>host</code>	The host field is deprecated. Use <code>host_uuid</code> instead.	Read only
<code>other-config</code>	Extra configuration information.	Read/write map parameter

pbd-create

```
1 xe pbd-create host-uuid=uuid_of_host sr-uuid=uuid_of_sr [device-config:  
key=corresponding_value]
```

Create a PBD on your XenServer host. The read-only `device-config` parameter can only be set on creation.

To add a mapping from 'path' to '/tmp', ensure that the command line contains the argument `device-config:path=/tmp`

For a full list of supported device-config key-value pairs on each SR type, see [Storage](#).

pbd-destroy

```
1 xe pbd-destroy uuid=uuid_of_pbd
```

Destroy the specified PBD.

pbd-plug

```
1 xe pbd-plug uuid=uuid_of_pbd
```

Attempts to plug in the PBD to the XenServer host. If this command succeeds, the referenced SR (and the VDIs contained within) becomes visible to the XenServer host.

pbd-unplug

```
1 xe pbd-unplug uuid=uuid_of_pbd
```

Attempt to unplug the PBD from the XenServer host.

PIF commands

Commands for working with PIFs (objects representing the physical network interfaces).

The PIF objects can be listed with the standard object listing command (`xe pif-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

PIF parameters

PIFs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the PIF	Read only
<code>device machine-readable</code>	Name of the interface (for example, eno12429np0)	Read only
<code>MAC</code>	The MAC address of the PIF	Read only
<code>other-config</code>	Extra PIF configuration <code>name: value</code> pairs.	Read/write map parameter
<code>physical</code>	If true, the PIF points to an actual physical network interface	Read only
<code>currently-attached</code>	Is the PIF currently attached on this host? true or false	Read only

Parameter Name	Description	Type
MTU	Maximum Transmission Unit of the PIF in bytes.	Read only
VLAN	VLAN tag for all traffic passing through this interface. -1 indicates that no VLAN tag is assigned	Read only
bond-master-of	The UUID of the bond this PIF is the main interface of (if any)	Read only
bond-slave-of	The UUID of the bond this PIF is part of (if any)	Read only
management	Is this PIF designated to be a management interface for the control domain	Read only
network-uuid	The unique identifier/object reference of the virtual network to which this PIF is connected	Read only
network-name-label	The name of the virtual network to which this PIF is connected	Read only
host-uuid	The unique identifier/object reference of the XenServer host to which this PIF is connected	Read only
host-name-label	The name of the XenServer host to which this PIF is connected	Read only
IP-configuration-mode	Type of network address configuration used; DHCP or static	Read only
IP	IP address of the PIF. Defined here when IP-configuration-mode is static; undefined when DHCP	Read only
netmask	Netmask of the PIF. Defined here when IP-configuration-mode is static; undefined when supplied by DHCP	Read only

Parameter Name	Description	Type
<code>gateway</code>	Gateway address of the PIF. Defined here when IP-configuration-mode is static; undefined when supplied by DHCP	Read only
<code>DNS</code>	DNS address of the PIF. Defined here when IP-configuration-mode is static; undefined when supplied by DHCP	Read only
<code>io_read_kbs</code>	Average read rate in kB/s for the device	Read only
<code>io_write_kbs</code>	Average write rate in kB/s for the device	Read only
<code>carrier</code>	Link state for this device	Read only
<code>vendor-id</code>	The ID assigned to NIC's vendor	Read only
<code>vendor-name</code>	The NIC vendor's name	Read only
<code>device-id</code>	The ID assigned by the vendor to this NIC model	Read only
<code>device-name</code>	The name assigned by the vendor to this NIC model	Read only
<code>speed</code>	Data transfer rate of the NIC	Read only
<code>duplex</code>	Duplexing mode of the NIC; full or half	Read only
<code>pci-bus-path</code>	PCI bus path address	Read only
<code>other-config:ethtoolspeed</code>	Sets the speed of connection in Mbps	Read/write
<code>other-config:ethtoolautoneg</code>	Set to no to disable autonegotiation of the physical interface or bridge. Default is yes.	Read/write
<code>other-config:ethtoolduplex</code>	Sets duplexing capability of the PIF, either full or half.	Read/write
<code>other-config:ethtool-rx</code>	Set to on to enable receive checksum, off to disable	Read/write

Parameter Name	Description	Type
<code>other-config:ethtool-tx</code>	Set to on to enable transmit checksum, off to disable	Read/write
<code>other-config:ethtool-sg</code>	Set to on to enable scatter gather, off to disable	Read/write
<code>other-config:ethtool-tso</code>	Set to on to enable TCP segmentation offload, off to disable	Read/write
<code>other-config:ethtool-ufo</code>	Set to on to enable UDP fragment offload, off to disable	Read/write
<code>other-config:ethtool-gso</code>	Set to on to enable generic segmentation offload, off to disable	Read/write
<code>other-config:domain</code>	Comma-separated list used to set the DNS search path	Read/write
<code>other-config:bondmiimon</code>	Interval between link liveness checks, in milliseconds	Read/write
<code>other-config:bonddowndelay</code>	Number of milliseconds to wait after link is lost before really considering the link to have gone. This parameter allows for transient link loss	Read/write
<code>other-config:bondupdelay</code>	Number of milliseconds to wait after the link comes up before really considering it up. Allows for links flapping up. Default is 31s to allow for time for switches to begin forwarding traffic.	Read/write
<code>disallow-unplug</code>	Value is true if this PIF is a dedicated storage NIC, false otherwise	Read/write

Note:

Changes made to the `other-config` fields of a PIF will only take effect after a reboot. Alternately, use the `xe pif-unplug` and `xe pif-plug` commands to cause the PIF configuration

to be rewritten.

pif-forget

```
1 xe pif-forget uuid=uuid_of_pif
```

Destroy the specified PIF object on a particular host.

pif-introduce

```
1 xe pif-introduce host-uuid=host_uuid mac=mac_address_for_pif device=
  interface_name
```

Create a PIF object representing a physical interface on the specified XenServer host.

pif-plug

```
1 xe pif-plug uuid=uuid_of_pif
```

Attempt to bring up the specified physical interface.

pif-reconfigure-ip

```
1 xe pif-reconfigure-ip uuid=uuid_of_pif [mode=dhcp|mode=static] gateway=
  network_gateway_address IP=static_ip_for_this_pif netmask=
  netmask_for_this_pif [DNS=dns_address]
```

Modify the IP address of the PIF. For static IP configuration, set the `mode` parameter to **static**, with the `gateway`, `IP`, and `netmask` parameters set to the appropriate values. To use DHCP, set the `mode` parameter to `DHCP` and leave the static parameters undefined.

Note:

Using static IP addresses on physical network interfaces connected to a port on a switch using Spanning Tree Protocol with STP Fast Link turned off (or unsupported) results in a period during which there is no traffic.

pif-reconfigure-ipv6

```
1 xe pif-reconfigure-ipv6 uuid=uuid_of_pif mode=mode [gateway=
  network_gateway_address] [IPv6=static_ip_for_this_pif] [DNS=
  dns_address]
```

Reconfigure the IPv6 address settings on a PIF.

pif-scan

```
1 xe pif-scan host-uuid=host_uuid
```

Scan for new physical interfaces on your XenServer host.

pif-set-primary-address-type

```
1 xe pif-set-primary-address-type uuid=uuid primary_address_type=
  address_type
```

Change the primary address type used by this PIF.

pif-unplug

```
1 xe pif-unplug uuid=uuid_of_pif
```

Attempt to bring down the specified physical interface.

Pool commands

Commands for working with pools. A *pool* is an aggregate of one or more XenServer hosts. A pool uses one or more shared storage repositories so that the VMs running on one host in the pool can be migrated in near-real time to another host in the pool. This migration happens while the VM is still running, without it needing to be shut down and brought back up.

Each XenServer host is really a pool consisting of a single member by default. When your XenServer host is joined to a pool, it is designated as a member. If the pool that the host is joined to consists of a single member, that member becomes the pool coordinator. If the pool that the host is joined to already has multiple members, one of these members is already the pool coordinator and remains so when the new host joins the pool.

The singleton pool object can be listed with the standard object listing command (`xe pool-list`). Its parameters can be manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

Pool parameters

Pools have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the pool	Read only
<code>name-label</code>	The name of the pool	Read/write
<code>name-description</code>	The description string of the pool	Read/write
<code>master</code>	The unique identifier/object reference of XenServer host designated as the pool's coordinator	Read only
<code>default-SR</code>	The unique identifier/object reference of the default SR for the pool	Read/write
<code>crash-dump-SR</code>	The unique identifier/object reference of the SR where any crash dumps for pool members are saved	Read/write
<code>metadata-vdis</code>	All known metadata VDIs for the pool	Read only
<code>suspend-image-SR</code>	The unique identifier/object reference of the SR where suspended VMs on pool members are saved	Read/write
<code>other-config</code>	A list of key-value pairs that specify extra configuration parameters for the pool	Read/write map parameter
<code>other-config: default_ha_timeout</code>	High availability timeout in seconds.	Read/write
<code>supported-sr-types</code>	SR types that this pool can use	Read only
<code>ha-enabled</code>	Value is true if HA is enabled for the pool, false otherwise	Read only
<code>ha-configuration</code>	Reserved for future use.	Read only
<code>ha-statefiles</code>	Lists the UUIDs of the VDIs being used by HA to determine storage health	Read only

Parameter Name	Description	Type
<code>ha-host-failures-to-tolerate</code>	The number of host failures to tolerate before sending a system alert	Read/write
<code>ha-plan-exists-for</code>	The number of hosts failures that can actually be handled, according to the calculations of the HA algorithm	Read only
<code>ha-allow-overcommit</code>	Value is true if the pool is allowed to be overcommitted, false otherwise	Read/write
<code>ha-overcommitted</code>	Value is true if the pool is overcommitted	Read only
<code>ha-reboot-vm-on-internal-shutdown</code>	Value is true if protected VMs are restarted on clean shut down, false otherwise.	Read/write
<code>blobs</code>	Binary data store	Read only
<code>live-patching-disabled</code>	Set to false to enable live patching. Set to true to disable live patching.	Read/write
<code>igmp-snooping-enabled</code>	Set to true to enable IGMP snooping. Set to false to disable IGMP snooping.	Read/write
<code>https-only</code>	Set to false to allow external clients that use the management API to connect to XenServer using either HTTPS over port 443 or HTTP over port 80. Set to true to block port 80 and require clients to exclusively connect using HTTPS over port 443.	Read/write

Parameter Name	Description	Type
<code>migration-compression</code>	Set to true to enable migration stream compression for your XenServer pool. Set to false to disable migration stream compression. The default is false. Can be overridden by <code>vm-migrate</code> command's <code>compress</code> parameter.	Read/write

pool-apply-edition

```
1 xe pool-apply-edition edition=cloud|standard-per-socket|premium-per-socket [uuid=uuid] [license-server-address=address] [license-server-port=port]
```

Apply an edition across the pool. Edition values include “cloud”, “standard-per-socket”, and “premium-per-socket”. In future, only “cloud” is supported.

pool-certificate-install

```
1 xe pool-certificate-install filename=file_name
```

Install an TLS certificate, pool-wide.

pool-certificate-list

```
1 xe pool-certificate-list
```

List all installed TLS certificates in a pool.

pool-certificate-sync

```
1 xe pool-certificate-sync
```

Sync TLS certificates and certificate revocation lists from pool coordinator to the other pool members.

pool-certificate-uninstall

```
1 xe pool-certificate-uninstall name=name
```

Uninstall a TLS certificate.

pool-crl-install

```
1 xe pool-crl-install filename=file_name
```

Install a TLS certificate revocation list, pool-wide.

pool-crl-list

```
1 xe pool-crl-list
```

List all installed TLS certificate revocation lists.

pool-crl-uninstall

```
1 xe pool-crl-uninstall name=name
```

Uninstall an TLS certificate revocation list.

pool-deconfigure-wlb

```
1 xe pool-deconfigure-wlb
```

Permanently remove the configuration for workload balancing.

pool-designate-new-master

```
1 xe pool-designate-new-master host-uuid=uuid_of_new_master
```

Instruct the specified member XenServer host to become the coordinator (formerly called “master”) of an existing pool. This command performs an orderly handover of the role of pool coordinator to another host in the resource pool. This command only works when the current pool coordinator is online. It is not a replacement for the emergency mode commands listed below.

pool-disable-external-auth

```
1 xe pool-disable-external-auth [uuid=uuid] [config=config]
```

Disables external authentication in all the hosts in a pool.

pool-disable-local-storage-caching

```
1 xe pool-disable-local-storage-caching uuid=uuid
```

Disable local storage caching across the pool.

pool-disable-redo-log

```
1 xe pool-disable-redo-log
```

Disable the redo log if in use, unless HA is enabled.

pool-dump-database

```
1 xe pool-dump-database file-name=filename_to_dump_database_into_(  
    on_client)
```

Download a copy of the entire pool database and dump it into a file on the client.

pool-enable-external-auth

```
1 xe pool-enable-external-auth auth-type=auth_type service-name=  
    service_name [uuid=uuid] [config:=config]
```

Enables external authentication in all the hosts in a pool. Note that some values of `auth-type` will require particular `config:` values.

pool-enable-local-storage-caching

```
1 xe pool-enable-local-storage-caching uuid=uuid
```

Enable local storage caching across the pool.

pool-enable-redo-log

```
1 xe pool-enable-redo-log sr-uuid=sr_uuid
```

Enable the redo log on the given SR if in use, unless HA is enabled.

pool-eject

```
1 xe pool-eject host-uuid=uuid_of_host_to_eject
```

Instruct the specified XenServer host to leave an existing pool.

pool-emergency-reset-master

```
1 xe pool-emergency-reset-master master-address=address_of_pool_master
```

Instruct a pool member host to reset its pool coordinator address to the new value and attempt to connect to it. Do not run this command on pool coordinators.

pool-emergency-transition-to-master

```
1 xe pool-emergency-transition-to-master
```

Instruct a member XenServer host to become the pool coordinator (formerly called “the pool master”). The XenServer host accepts this command only after the host has transitioned to emergency mode. Emergency mode means it is a member of a pool whose coordinator has disappeared from the network and cannot be contacted after some number of retries.

If the host password has been modified since the host joined the pool, this command can cause the password of the host to reset. For more information, see ([User commands](#)).

pool-ha-enable

```
1 xe pool-ha-enable heartbeat-sr-uuids=uuid_of_heartbeat_sr
```

Enable high availability on the resource pool, using the specified SR UUID as the central storage heartbeat repository.

pool-ha-disable

```
1 xe pool-ha-disable
```

Disables the high availability feature on the resource pool.

pool-ha-compute-hypothetical-max-host-failures-to-tolerate

Compute the maximum number of host failures to tolerate under the current pool configuration.

pool-ha-compute-max-host-failures-to-tolerate

```
1 xe pool-ha-compute-hypothetical-max-host-failures-to-tolerate [vm-uuid=  
vm_uuid] [restart-priority=restart_priority]
```

Compute the maximum number of host failures to tolerate with the supplied, proposed protected VMs.

pool-initialize-wlb

```
1 xe pool-initialize-wlb wlb_url=url wlb_username=wb_username  
wlb_password=wlb_password xenserver_username=username  
xenserver_password=password
```

Initialize workload balancing for the current pool with the target Workload Balancing server.

pool-join

```
1 xe pool-join master-address=address master-username=username master-  
password=password
```

Instruct your XenServer host to join an existing pool.

pool-management-reconfigure

```
1 xe pool-management-reconfigure [network-uuid=network-uuid]
```

Reconfigures the management interface of all the hosts in the pool to use the specified network interface, which is the interface that is used to connect to the XenCenter. The command rewrites the MANAGEMENT_INTERFACE key in `/etc/xensource-inventory` for all the hosts in the pool.

If the device name of an interface (which must have an IP address) is specified, the XenServer pool coordinator immediately rebinds. This command works both in normal and emergency mode.

From the network UUID specified, UUID of the PIF object is identified and mapped to the XenServer host, which determines which IP address to rebind to itself. It must not be in emergency mode when this command is run.

Warning:

Be careful when using this CLI command off-host and ensure that you have network connectivity on the new interface. Use `xe pif-reconfigure` to set one up first. Otherwise, subsequent CLI commands are unable to reach the XenServer host.

pool-recover-slaves

```
1 xe pool-recover-slaves
```

Instruct the pool coordinator to try to reset the address of all members currently running in emergency mode. This command is typically used after `pool-emergency-transition-to-master` has been used to set one of the members as the new pool coordinator.

pool-restore-database

```
1 xe pool-restore-database file-name=filename_to_restore_from_on_client [
    dry-run=true|false]
```

Upload a database backup (created with `pool-dump-database`) to a pool. On receiving the upload, the pool coordinator restarts itself with the new database.

There is also a *dry run* option, which allows you to check that the pool database can be restored without actually perform the operation. By default, `dry-run` is set to false.

pool-retrieve-wlb-configuration

```
1 xe pool-retrieve-wlb-configuration
```

Retrieves the pool optimization criteria from the Workload Balancing server.

pool-retrieve-wlb-diagnostics

```
1 xe pool-retrieve-wlb-diagnostics [filename=file_name]
```

Retrieves diagnostics from the Workload Balancing server.

pool-retrieve-wlb-recommendations

```
1 xe pool-retrieve-wlb-recommendations
```

Retrieves VM migrate recommendations for the pool from the Workload Balancing server.

pool-retrieve-wlb-report

```
1 xe pool-retrieve-wlb-report report=report [filename=file_name]
```

Retrieves reports from the Workload Balancing server.

pool-secret-rotate

```
1 xe pool-secret-rotate
```

Rotate the pool secret.

The pool secret is a secret shared among the XenServer hosts in a pool that enables the host to prove its membership to a pool. Users with the Pool Admin role can view this secret when connecting to the host over SSH. Rotate the pool secret if one of these users leaves your organization or loses their Pool Admin role.

pool-send-test-post

```
1 xe pool-send-test-post dest-host=destination_host dest-port=
  destination_port body=post_body
```

Send the given body to the given host and port, using HTTPS, and print the response. This is used for debugging the TLS layer.

pool-send-wlb-configuration

```
1 xe pool-send-wlb-configuration [config=config]
```

Sets the pool optimization criteria for the Workload Balancing server.

pool-sync-database

```
1 xe pool-sync-database
```

Force the pool database to be synchronized across all hosts in the resource pool. This command is not necessary in normal operation since the database is regularly automatically replicated. However, the command can be useful for ensuring changes are rapidly replicated after performing a significant set of CLI operations.

Set https-only

```
1 xe pool-param-set [uuid=pool-uuid] [https-only=true | false]
```

Enables or disables the blocking of port 80 on the management interface of XenServer hosts.

PVS Accelerator commands

Commands for working with the PVS Accelerator.

pvs-cache-storage-create

```
1 xe pvs-cache-storage-create sr-uuid=sr_uuid pvs-site-uuid=pvs_site_uuid  
   size=size
```

Configure a PVS cache on a given SR for a given host.

pvs-cache-storage-destroy

```
1 xe pvs-cache-storage-destroy uuid=uuid
```

Remove a PVS cache.

pvs-proxy-create

```
1 xe pvs-proxy-create pvs-site-uuid=pvs_site_uuid vif-uuid=vif_uuid
```

Configure a VM/VIF to use a PVS proxy.

pvs-proxy-destroy

```
1 xe pvs-proxy-destroy uuid=uuid
```

Remove (or switch off) a PVS proxy for this VIF/VM.

pvs-server-forget

```
1 xe pvs-server-forget uuid=uuid
```

Forget a PVS server.

pvs-server-introduce

```
1 xe pvs-server-introduce addresses=addresses first-port=first_port last-port=last_port pvs-site-uuid=pvs_site_uuid
```

Introduce new PVS server.

pvs-site-forget

```
1 xe pvs-site-forget uuid=uuid
```

Forget a PVS site.

pvs-site-introduce

```
1 xe pvs-site-introduce name-label=name_label [name-description=name_description] [pvs-uuid=pvs_uuid]
```

Introduce new PVS site.

Storage Manager commands

Commands for controlling Storage Manager plug-ins.

The storage manager objects can be listed with the standard object listing command (`xe sm-list`). The parameters can be manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

SM parameters

SMs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the SM plug-in	Read only
<code>name-label</code>	The name of the SM plug-in	Read only
<code>name-description</code>	The description string of the SM plug-in	Read only
<code>type</code>	The SR type that this plug-in connects to	Read only

Parameter Name	Description	Type
<code>vendor</code>	Name of the vendor who created this plug-in	Read only
<code>copyright</code>	Copyright statement for this SM plug-in	Read only
<code>required-api-version</code>	Minimum SM API version required on the XenServer host	Read only
<code>configuration</code>	Names and descriptions of device configuration keys	Read only
<code>capabilities</code>	Capabilities of the SM plug-in	Read only
<code>driver-filename</code>	The file name of the SR driver.	Read only

Manage SSH Access on Hosts and Pools

Overview

In XenServer 8.4, SSH access to dom0 is now enabled by default. This release introduces the capability to programmatically control SSH access using either the XenAPI or the xe command-line interface (CLI).

This enhancement allows administrators to manage SSH access at both individual host and entire pool levels, eliminating the need for manual interaction with each server console.

Benefits

- **Automated Management:** Enable or disable SSH on individual hosts or across an entire pool programmatically.
- **Reduced Administrative Overhead:** Significantly reduce manual administrative effort required for managing SSH access.
- **Enhanced Workflows:** Improve automation scripts and bolster security workflows by integrating SSH control into existing management systems.

Compatibility and Caveats

- **No Additional Requirements:** This feature does not require any additional configuration or specific compatibility adjustments.

- **Service Behavior:** Disabling SSH will stop and disable the sshd service on the specified host(s).
- **Existing Sessions:** Any existing SSH sessions will remain active and uninterrupted even after disabling the service.
- **New Connections:** New connection attempts will be refused once SSH is disabled.
- **Pool-wide Operations:** Operations performed across an entire pool will continue even if one or more hosts within the pool encounter a failure.

A summary message will be provided, detailing any failed operations.

CLI Configuration

Enable SSH on a host

```
1 xe host-enable-ssh host=<host_uuid_or_name_label>
```

This command starts and enables the `sshd` service on the specified host.

For valid `host` selectors, see [Host Selectors](#)

.

Disable SSH on a host

```
1 xe host-disable-ssh host=<host_uuid_or_name_label>
```

This command stops and disables the sshd service on the specified host.

Enable SSH on all hosts in a pool

```
1 xe pool-enable-ssh
```

This command starts and enables the sshd service on all hosts in the current pool.

Disable SSH on all hosts in a pool

```
1 xe pool-disable-ssh
```

This command stops and disables the sshd service on all hosts in the current pool.

Example Output (Pool-level Operation)

```
1 xe pool-enable-ssh
2
3 Some of hosts failed to enable SSH access.
4 hosts: d734ff47-f429-4bd4-93f7-245cffbdc79e (genus-35-003d)
```

If a host fails, the command continues processing the remaining hosts and reports failures at the end.

Snapshot commands

Commands for working with snapshots.

snapshot-clone

```
1 xe snapshot-clone new-name-label=name_label [uuid=uuid] [new-name-  
description=description]
```

Create a new template by cloning an existing snapshot, using storage-level fast disk clone operation where available.

snapshot-copy

```
1 xe snapshot-copy new-name-label=name_label [uuid=uuid] [new-name-  
description=name_description] [sr-uuid=sr_uuid]
```

Create a new template by copying an existing VM, but without using storage-level fast disk clone operation (even if this is available). The disk images of the copied VM are guaranteed to be ‘full images’ - i.e. not part of a CoW chain.

snapshot-destroy

```
1 xe snapshot-destroy [uuid=uuid] [snapshot-uuid=snapshot_uuid]
```

Destroy a snapshot. This leaves the storage associated with the snapshot intact. To delete storage too, use snapshot-uninstall.

snapshot-disk-list

```
1 xe snapshot-disk-list [uuid=uuid] [snapshot-uuid=snapshot_uuid] [vbd-  
params=vbd_params] [vdi-params=vdi_params]
```

List the disks on the selected VM(s).

snapshot-export-to-template

```
1 xe snapshot-export-to-template filename=file_name snapshot-uuid=  
snapshot_uuid [preserve-power-state=true|false]
```

Export a snapshot to *file name*.

snapshot-reset-powerstate

```
1 xe snapshot-reset-powerstate [uuid=uuid] [snapshot-uuid=snapshot_uuid]
   [--force]
```

Force the VM power state to halted in the management toolstack database only. This command is used to recover a snapshot that is marked as ‘suspended’. This is a potentially dangerous operation: you must ensure that you do not need the memory image anymore. You will not be able to resume your snapshot anymore.

snapshot-revert

```
1 xe snapshot-revert [uuid=uuid] [snapshot-uuid=snapshot_uuid]
```

Revert an existing VM to a previous checkpointed or snapshot state.

snapshot-uninstall

```
1 xe snapshot-uninstall [uuid=uuid] [snapshot-uuid=snapshot_uuid] [--
   force]
```

Uninstall a snapshot. This operation will destroy those VDIs that are marked RW and connected to this snapshot only. To simply destroy the VM record, use `snapshot-destroy`.

SR commands

Commands for controlling SRs (storage repositories).

The SR objects can be listed with the standard object listing command (`xe sr-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

SR parameters

SRs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the SR	Read only
<code>name-label</code>	The name of the SR	Read/write
<code>name-description</code>	The description string of the SR	Read/write

Parameter Name	Description	Type
<code>host</code>	The storage repository host name	Read only
<code>allowed-operations</code>	List of the operations allowed on the SR in this state	Read only set parameter
<code>current-operations</code>	List of the operations that are currently in progress on this SR	Read only set parameter
<code>VDIs</code>	Unique identifier/object reference for the virtual disks in this SR	Read only set parameter
<code>PBDs</code>	Unique identifier/object reference for the PBDs attached to this SR	Read only set parameter
<code>virtual-allocation</code>	Sum of virtual-size values of all VDIs in this storage repository (in bytes)	Read only
<code>physical-utilisation</code>	Physical space currently utilized on this SR, in bytes. For thin provisioned disk formats, physical utilization may be less than virtual allocation.	Read only
<code>physical-size</code>	Total physical size of the SR, in bytes	Read only
<code>type</code>	Type of the SR, used to specify the SR back-end driver to use	Read only
<code>content-type</code>	The type of the SR's content. Used to distinguish ISO libraries from other SRs. For storage repositories that store a library of ISOs, the <code>content-type</code> must be set to <code>iso</code> . In other cases, we recommend that you set this parameter either to empty, or the string user.	Read only
<code>shared</code>	Value is true if this SR can be shared between multiple hosts, false otherwise.	Read/write

Parameter Name	Description	Type
<code>introduced-by</code>	The <code>drtask</code> (if any) which introduced the SR	Read only
<code>is-tools-sr</code>	Value is true if this is the SR that contains the Tools ISO VDIs. false otherwise.	Read only
<code>other-config</code>	List of key-value pairs that specify extra configuration parameters for the SR	Read/write map parameter
<code>sm-config</code>	SM dependent data	Read only map parameter
<code>blobs</code>	Binary data store	Read only
<code>local-cache-enabled</code>	Value is true if this SR is assigned to be the local cache for its host, false otherwise.	Read only
<code>tags</code>	User-specified tags for categorization purposes	Read/write set parameter
<code>clustered</code>	Value is true if the SR is using aggregated local storage, false otherwise.	Read only

sr-create

```
1 xe sr-create name=label=name physical-size=size type=type content-type=
  content_type device-config:config_name=value [host-uuid=host_uuid] [
  shared=true|false]
```

Creates an SR on the disk, introduces it into the database, and creates a PBD attaching the SR to the XenServer host. If `shared` is set to **true**, a PBD is created for each XenServer host in the pool. If `shared` is not specified or set to **false**, a PBD is created only for the XenServer host specified with `host-uuid`.

The exact `device-config` parameters differ depending on the device `type`. For details of these parameters across the different storage back-ends, see [Create an SR](#).

sr-data-source-forget

```
1 xe sr-data-source-forget data-source=data_source
```

Stop recording the specified data source for a SR, and forget all of the recorded data.

sr-data-source-list

```
1 xe sr-data-source-list
```

List the data sources that can be recorded for a SR.

sr-data-source-query

```
1 xe sr-data-source-query data-source=data_source
```

Query the last value read from a SR data source.

sr-data-source-record

```
1 xe sr-data-source-record data-source=data_source
```

Record the specified data source for a SR.

sr-destroy

```
1 xe sr-destroy uuid=sr_uuid
```

Destroys the specified SR on the XenServer host.

sr-enable-database-replication

```
1 xe sr-enable-database-replication uuid=sr_uuid
```

Enables XAPI database replication to the specified (shared) SR.

sr-disable-database-replication

```
1 xe sr-disable-database-replication uuid=sr_uuid
```

Disables XAPI database replication to the specified SR.

sr-forget

```
1 xe sr-forget uuid=sr_uuid
```

The XAPI agent forgets about a specified SR on the XenServer host. When the XAPI agent forgets an SR, the SR is detached and you cannot access VDIs on it, but it remains intact on the source media (the data is not lost).

sr-introduce

```
1 xe sr-introduce name=label=name physical-size=physical_size type=type  
  content-type=content_type uuid=sr_uuid
```

Just places an SR record into the database. Use `device-config` to specify additional parameters in the form `device-config:parameter_key=parameter_value`, for example:

```
1 xe sr-introduce device-config:device=/dev/sdb1
```

Note:

This command is never used in normal operation. This advanced operation might be useful when an SR must be reconfigured as shared after it was created or to help recover from various failure scenarios.

sr-probe

```
1 xe sr-probe type=type [host-uuid=host_uuid] [device-config:config_name=  
  value]
```

Performs a scan of the backend, using the provided `device-config` keys. If the `device-config` is complete for the SR back-end, this command returns a list of the SRs present on the device, if any. If the `device-config` parameters are only partial, a back-end-specific scan is performed, returning results that guide you in improving the remaining `device-config` parameters. The scan results are returned as XML specific to the back end, printed on the CLI.

The exact `device-config` parameters differ depending on the device `type`. For details of these parameters across the different storage back-ends, see [Storage](#).

sr-probe-ext

```
1 xe sr-probe-ext type=type [host-uuid=host_uuid] [device-config:=config]  
  [sm-config:-sm_config]
```

Perform a storage probe. The `device-config` parameters can be specified by for example `device-config:devs=/dev/sdb1`. Unlike `sr-probe`, this command returns results in the same human-readable format for every SR type.

sr-scan

```
1 xe sr-scan uuid=sr_uuid
```

Force an SR scan, syncing the XAPI database with VDIs present in the underlying storage substrate.

sr-update

```
1 xe sr-update uuid=uuid
```

Refresh the fields of the SR object in the database.

lvhd-enable-thin-provisioning

```
1 xe lvhd-enable-thin-provisioning sr-uuid=sr_uuid initial-allocation=  
initial_allocation allocation-quantum=allocation_quantum
```

Enable thin-provisioning on an LVHD SR.

Subject commands

Commands for working with subjects.

session-subject-identifier-list

```
1 xe session-subject-identifier-list
```

Return a list of all the user subject ids of all externally-authenticated existing sessions.

session-subject-identifier-logout

```
1 xe session-subject-identifier-logout subject-identifier=  
subject_identifier
```

Log out all externally-authenticated sessions associated to a user subject id.

session-subject-identifier-logout-all

```
1 xe session-subject-identifier-logout-all
```

Log out all externally-authenticated sessions.

subject-add

```
1 xe subject-add subject-name=subject_name
```

Add a subject to the list of subjects that can access the pool.

subject-remove

```
1 xe subject-remove subject-uuid=subject_uuid
```

Remove a subject from the list of subjects that can access the pool.

subject-role-add

```
1 xe subject-role-add uuid=uuid [role-name=role_name] [role-uuid=role_uuid]
```

Add a role to a subject.

subject-role-remove

```
1 xe subject-role-remove uuid=uuid [role-name=role_name] [role-uuid=role_uuid]
```

Remove a role from a subject.

secret-create

```
1 xe secret-create value=value
```

Create a secret.

secret-destroy

```
1 xe secret-destroy uuid=uuid
```

Destroy a secret.

Task commands

Commands for working with long-running asynchronous tasks. These commands are tasks such as starting, stopping, and suspending a virtual machine. The tasks are typically made up of a set of other atomic subtasks that together accomplish the requested operation.

The task objects can be listed with the standard object listing command (`xe task-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

Task parameters

Tasks have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the Task	Read only
<code>name-label</code>	The name of the Task	Read only
<code>name-description</code>	The description string of the Task	Read only
<code>resident-on</code>	The unique identifier/object reference of the host on which the task is running	Read only
<code>status</code>	Status of the Task	Read only
<code>progress</code>	If the Task is still pending, this field contains the estimated percentage complete, from 0 to 1. If the Task has completed, successfully or unsuccessfully, the value is 1.	Read only
<code>type</code>	If the Task has successfully completed, this parameter contains the type of the encoded result. The type is the name of the class whose reference is in the result field. Otherwise, this parameter's value is undefined	Read only

Parameter Name	Description	Type
<code>result</code>	If the Task has completed successfully, this field contains the result value, either Void or an object reference; otherwise, this parameter's value is undefined	Read only
<code>error_info</code>	If the Task has failed, this parameter contains the set of associated error strings. Otherwise, this parameter's value is undefined	Read only
<code>allowed_operations</code>	List of the operations allowed in this state	Read only
<code>created</code>	Time the task has been created	Read only
<code>finished</code>	Time task finished (that is, succeeded or failed). If task-status is pending, then the value of this field has no meaning	Read only
<code>subtask_of</code>	Contains the UUID of the tasks this task is a subtask of	Read only
<code>subtasks</code>	Contains the UUIDs of all the subtasks of this task	Read only

task-cancel

```
1 xe task-cancel [uuid=task_uuid]
```

Direct the specified Task to cancel and return.

Template commands

Commands for working with VM templates.

Templates are essentially VMs with the `is-a-template` parameter set to **true**. A template is a “gold image” that contains all the various configuration settings to instantiate a specific VM. XenServer ships with a base set of templates, which are generic “raw” VMs that can boot an OS vendor installation CD (for example: RHEL, CentOS, SLES, Windows). You can create VMs, configure them in standard

forms for your particular needs, and save a copy of them as templates for future use in VM deployment.

The template objects can be listed with the standard object listing command (`xe template-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

Note:

Templates cannot be directly converted into VMs by setting the `is-a-template` parameter to **false**. Setting `is-a-template` parameter to **false** is not supported and results in a VM that cannot be started.

VM template parameters

Templates have the following parameters:

- `uuid` (read only) the unique identifier/object reference for the template
- `name-label` (read/write) the name of the template
- `name-description` (read/write) the description string of the template
- `user-version` (read/write) string for creators of VMs and templates to put version information
- `is-a-template` (read/write) true if this VM is a template.
Template VMs can never be started, they are used only for cloning other VMs. After this value has been set to true, it cannot be reset to false. Template VMs cannot be converted into VMs using this parameter.

You can convert a VM to a template with:

```
1 xe vm-param-set uuid=<vm uuid> is-a-template=true
```

- `is-control-domain` (read only) true if this is a control domain (domain 0 or a driver domain)
- `power-state` (read only) current power state. The value is always halted for a template
- `memory-dynamic-max` (read only) dynamic maximum memory in bytes.
Currently unused, but if changed the following constraint must be obeyed: `memory_static_max >= memory_dynamic_max >= memory_dynamic_min >= memory_static_min`.
- `memory-dynamic-min` (read/write) dynamic minimum memory in bytes.
Currently unused, but if changed the same constraints for `memory-dynamic-max` must be obeyed.

- `memory-static-max` (read/write) statically set (absolute) maximum memory in bytes. This field is the main value used to determine the amount of memory assigned to a VM.
- `memory-static-min` (read/write) statically set (absolute) minimum memory in bytes. This field represents the absolute minimum memory, and `memory-static-min` must be less than `memory-static-max`.
This value is unused in normal operation, but the previous constraint must be obeyed.
- `suspend-VDI-uuid` (read only) the VDI that a suspend image is stored on (has no meaning for a template)
- `VCPUs-params` (read/write map parameter) configuration parameters for the selected vCPU policy.
- `VCPUs-max` (read/write) maximum number of vCPUs
- `VCPUs-at-startup` (read/write) boot number of vCPUs
- `actions-after-crash` (read/write) action to take when a VM based on this template crashes
- `console-uuids` (read only set parameter) virtual console devices
- `platform` (read/write map parameter) platform specific configuration

To disable the emulation of a parallel port for guests:

```
1 xe vm-param-set uuid=<vm_uuid> platform:parallel=none
```

To disable the emulation of a serial port:

```
1 xe vm-param-set uuid=<vm_uuid> platform:hvm_serial=none
```

To disable the emulation of a USB controller and a USB tablet device:

```
1 xe vm-param-set uuid=<vm_uuid> platform:usb=false
2 xe vm-param-set uuid=<vm_uuid> platform:usb_tablet=false
```

- `allowed-operations` (read only set parameter) list of the operations allowed in this state
- `current-operations` (read only set parameter) list of the operations that are currently in progress on this template
- `allowed-VBD-devices` (read only set parameter) list of VBD identifiers available for use, represented by integers of the range 0–254. This list is informational only, and other devices may be used (but may not work).
- `allowed-VIF-devices` (read only set parameter) list of VIF identifiers available for use, represented by integers of the range 0–6. This list is informational only, and other devices may be used (but may not work).

- `HVM-boot-policy` (read/write) the boot policy for guests. Either BIOS Order or an empty string.
- `HVM-boot-params` (read/write map parameter) the order key controls the guest boot order, represented as a string where each character is a boot method: d for the CD/DVD, c for the root disk, and n for network PXE boot. The default is dc.
- `PV-kernel` (read/write) path to the kernel
- `PV-ramdisk` (read/write) path to the `initrd`
- `PV-args` (read/write) string of kernel command line arguments
- `PV-legacy-args` (read/write) string of arguments to make legacy VMs based on this template boot
- `PV-bootloader` (read/write) name of or path to bootloader
- `PV-bootloader-args` (read/write) string of miscellaneous arguments for the bootloader
- `last-boot-CPU-flags` (read only) describes the CPU flags on which a VM based on this template was last booted; not populated for a template
- `resident-on` (read only) the XenServer host on which a VM based on this template is resident. Appears as `not in database` for a template
- `affinity` (read/write) the XenServer host which a VM based on this template has preference for running on. Used by the `xe vm-start` command to decide where to run the VM. Also known as the home server.
- `groups` (read/write) the VM group this VM belongs to. A VM group defines where XenServer attempts to start the VM or evacuate it to.
- `other-config` (read/write map parameter) list of key-value pairs that specify extra configuration parameters for the template
- `start-time` (read only) timestamp of the date and time that the metrics for a VM based on this template were read, in the form `yyyymmddThh:mm:ss z`, where z is the single-letter military timezone indicator, for example, Z for UTC(GMT). Set to 1 Jan 1970 Z (beginning of Unix/POSIX epoch) for a template
- `install-time` (read only) timestamp of the date and time that the metrics for a VM based on this template were read, in the form `yyyymmddThh:mm:ss z`, where z is the single-letter military timezone indicator, for example, Z for UTC (GMT). Set to 1 Jan 1970 Z (beginning of Unix/POSIX epoch) for a template
- `memory-actual` (read only) the actual memory being used by a VM based on this template; 0 for a template

- **VCPUs-number** (read only) the number of virtual CPUs assigned to a VM based on this template; 0 for a template
- **VCPUs-Utilization** (read only map parameter) list of virtual CPUs and their weight read only map parameter **os-version** the version of the operating system for a VM based on this template. Appears as **not in database** for a template
- **PV-drivers-version** (read only map parameter) the versions of the paravirtualized drivers for a VM based on this template. Appears as **not in database** for a template
- **PV-drivers-detected** (read only) flag for latest version of the paravirtualized drivers for a VM based on this template. Appears as **not in database** for a template
- **memory** (read only map parameter) memory metrics reported by the agent on a VM based on this template. Appears as **not in database** for a template
- **disks** (read only map parameter) disk metrics reported by the agent on a VM based on this template. Appears as **not in database** for a template
- **networks** (read only map parameter) network metrics reported by the agent on a VM based on this template. Appears as **not in database** for a template
- **other** (read only map parameter) other metrics reported by the agent on a VM based on this template. Appears as **not in database** for a template
- **guest-metrics-last-updated** (read only) timestamp when the in-guest agent performed the last write to these fields. In the form **yyyymmddThh:mm:ss z**, where z is the single-letter military timezone indicator, for example, Z for UTC (GMT)
- **actions-after-shutdown** (read/write) action to take after the VM has shutdown
- **actions-after-reboot** (read/write) action to take after the VM has rebooted
- **possible-hosts** (read only) list of hosts that can potentially host the VM
- **HVM-shadow-multiplier** (read/write) multiplier applied to the amount of shadow that is made available to the guest
- **dom-id** (read only) domain ID (if available, -1 otherwise)
- **recommendations** (read only) XML specification of recommended values and ranges for properties of this VM
- **xenstore-data** (read/write map parameter) data to be inserted into the **xenstore** tree (**/local/domain/*domid*/vmdata**) after the VM is created.
- **is-a-snapshot** (read only) Value is true if this template is a VM snapshot
- **snapshot_of** (read only) the UUID of the VM that this template is a snapshot of
- **snapshots** (read only) the UUIDs of any snapshots that have been taken of this template

- `snapshot_time` (read only) the timestamp of the most recent VM snapshot taken
- `memory-target` (read only) the target amount of memory set for this template
- `blocked-operations` (read/write map parameter) lists the operations that cannot be performed on this template
- `last-boot-record` (read only) record of the last boot parameters for this template, in XML format
- `ha-always-run` (read/write) Value is true if an instance of this template is always restarted on another host if there is a failure of the host it is resident on. This parameter is now deprecated. Use the `ha-restartpriority` parameter instead.
- `ha-restart-priority` (read only) restart or best-effort read/write blobs binary data store
- `live` (read only) relevant only to a running VM.

template-export

```
1 xe template-export template-uuid=uuid_of_existing_template filename=
  filename_for_new_template
```

Exports a copy of a specified template to a file with the specified new file name.

template-uninstall

```
1 xe template-uninstall template-uuid=template_uuid [--force]
```

Uninstall a custom template. This operation will destroy those VDIs that are marked as ‘owned’ by this template.

Update commands

The following section contains XenServer host update commands.

The update objects can be listed with the standard object listing command (`xe update-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

Update parameters

XenServer host updates have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the update	Read only
<code>host</code>	The list of hosts that this update is applied to	Read only
<code>host-uuid</code>	The unique identifier for the XenServer host to query	Read only
<code>name-label</code>	The name of the update	Read only
<code>name-description</code>	The description string of the update	Read only
<code>applied</code>	Whether or not the update has been applied; true or false	Read only
<code>installation-size</code>	The size of the update in bytes	Read only
<code>after-apply-guidance</code>	Whether the XAPI toolstack or the host requires a restart	Read only
<code>version</code>	The version of the update	Read only

update-upload

```
1 xe update-upload file-name=update_filename
```

Upload a specified update file to the XenServer host. This command prepares an update to be applied. On success, the UUID of the uploaded update is printed. If the update has previously been uploaded, `UPDATE_ALREADY_EXISTS` error is returned instead and the patch is not uploaded again.

update-precheck

```
1 xe update-precheck uuid=update_uuid host-uuid=host_uuid
```

Run the pre-checks contained within the specified update on the specified XenServer host.

update-destroy

```
1 xe update-destroy uuid=update_file_uuid
```

Deletes an update file that has not been applied from the pool. Can be used to delete an update file that cannot be applied to the hosts.

update-apply

```
1 xe update-apply host-uuid=host_uuid uuid=update_file_uuid
```

Apply the specified update file.

update-pool-apply

```
1 xe update-pool-apply uuid=update_uuid
```

Apply the specified update to all XenServer hosts in the pool.

update-introduce

```
1 xe update-introduce vdi-uuid=vdi_uuid
```

Introduce update VDI.

update-pool-clean

```
1 xe update-pool-clean uuid=uuid
```

Removes the update's files from all hosts in the pool.

User commands**user-password-change**

```
1 xe user-password-change old=old_password new=new_password
```

Changes the password of the logged-in user. The old password field is not checked because you require supervisor privilege to use this command.

VBD commands

Commands for working with VBDs (Virtual Block Devices).

A VBD is a software object that connects a VM to the VDI, which represents the contents of the virtual disk. The VBD has the attributes which tie the VDI to the VM (is it bootable, its read/write metrics, and so on). The VDI has the information on the physical attributes of the virtual disk (which type of SR, whether the disk is sharable, whether the media is read/write or read only, and so on).

The VBD objects can be listed with the standard object listing command (`xe vbd-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

VBD parameters

VBDs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the VBD	Read only
<code>vm-uuid</code>	The unique identifier/object reference for the VM this VBD is attached to	Read only
<code>vm-name-label</code>	The name of the VM this VBD is attached to	Read only
<code>vdi-uuid</code>	The unique identifier/object reference for the VDI this VBD is mapped to	Read only
<code>vdi-name-label</code>	The name of the VDI this VBD is mapped to	Read only
<code>empty</code>	If true , this VBD represents an empty drive	Read only
<code>device</code>	The device seen by the guest, for example <code>hda</code>	Read only
<code>userdevice</code>	Device number specified by the device parameter during <code>vbd-create</code> , for example, 0 for <code>hda</code> , 1 for <code>hdb</code> , and so on	Read/write
<code>bootable</code>	Value is true if this VBD is bootable	Read/write
<code>mode</code>	The mode the VBD is mounted with	Read/write
<code>type</code>	How the VBD appears to the VM, for example disk or CD	Read/write
<code>currently-attached</code>	Value is true if the VBD is attached on this host, false otherwise	Read only

Parameter Name	Description	Type
<code>storage-lock</code>	Value is true if a storage-level lock was acquired	Read only
<code>status-code</code>	Error/success code associated with the last attach operation	Read only
<code>status-detail</code>	Error/success information associated with the last attach operation status	Read only
<code>qos_algorithm_type</code>	The prioritization algorithm to use	Read/write
<code>qos_algorithm_params</code>	Parameters for the chosen prioritization algorithm	Read/write map parameter
<code>qos_supported_algorithms</code>	Supported prioritization algorithms for this VBD	Read only set parameter
<code>io_read_kbs</code>	Average read rate in kB per second for this VBD	Read only
<code>io_write_kbs</code>	Average write rate in kB per second for this VBD	Read only
<code>allowed-operations</code>	List of the operations allowed in this state. This list is advisory only and the server state may have changed by the time this field is read by a client.	Read only set parameter
<code>current-operations</code>	Links each of the running tasks using this object (by reference) to a <code>current_operation</code> enum which describes the nature of the task.	Read only set parameter
<code>unpluggable</code>	Value is true if this VBD supports hot unplug	Read/write
<code>attachable</code>	Value is true if the device can be attached	Read only
<code>other-config</code>	Extra configuration	Read/write map parameter

vbd-create

```
1 xe vbd-create vm-uuid=uuid_of_the_vm device=device_value vdi-uuid=
   uuid_of_vdi_to_connect_to [bootable=true] [type=Disk|CD] [mode=RW|RO]
```

```
]
```

Create a VBD on a VM.

The allowable values for the `device` field are integers 0–254, and the number must be unique for each VM. The current allowable values can be seen in the `allowed-vbd-devices` parameter on the specified VM. This is seen as `userdevice` in the `vbd` parameters.

If the `type` is `Disk`, `vdi-uuid` is required. Mode can be `RO` or `RW` for a Disk.

If the `type` is `CD`, `vdi-uuid` is optional. If no VDI is specified, an empty VBD is created for the CD. Mode must be `RO` for a CD.

vbd-destroy

```
1 xe vbd-destroy uuid=uuid_of_vbd
```

Destroy the specified VBD.

If the VBD has its `other-config:owner` parameter set to `true`, the associated VDI is also destroyed.

vbd-eject

```
1 xe vbd-eject uuid=uuid_of_vbd
```

Remove the media from the drive represented by a VBD. This command only works if the media is of a removable type (a physical CD or an ISO). Otherwise, an error message `VBD_NOT_REMOVABLE_MEDIA` is returned.

vbd-insert

```
1 xe vbd-insert uuid=uuid_of_vbd vdi-uuid=uuid_of_vdi_containing_media
```

Insert new media into the drive represented by a VBD. This command only works if the media is of a removable type (a physical CD or an ISO). Otherwise, an error message `VBD_NOT_REMOVABLE_MEDIA` is returned.

vbd-plug

```
1 xe vbd-plug uuid=uuid_of_vbd
```

Attempt to attach the VBD while the VM is in the running state.

vbd-unplug

```
1 xe vbd-unplug uuid=uuid_of_vbd
```

Attempts to detach the VBD from the VM while it is in the running state.

VDI commands

Commands for working with VDIs (Virtual Disk Images).

A VDI is a software object that represents the contents of the virtual disk seen by a VM. This is different to the VBD, which is an object that ties a VM to the VDI. The VDI has the information on the physical attributes of the virtual disk (which type of SR, whether the disk is sharable, whether the media is read/write or read only, and so on). The VBD has the attributes that tie the VDI to the VM (is it bootable, its read/write metrics, and so on).

The VDI objects can be listed with the standard object listing command (`xe vdi-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

VDI parameters

VDIs have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the VDI	Read only
<code>name-label</code>	The name of the VDI	Read/write
<code>name-description</code>	The description string of the VDI	Read/write
<code>allowed-operations</code>	A list of the operations allowed in this state	Read only set parameter
<code>current-operations</code>	A list of the operations that are currently in progress on this VDI	Read only set parameter
<code>sr-uuid</code>	SR in which the VDI resides	Read only
<code>vbd-uuids</code>	A list of VBDs that refer to this VDI	Read only set parameter
<code>crashdump-uuids</code>	List of crash dumps that refer to this VDI	Read only set parameter

Parameter Name	Description	Type
<code>virtual-size</code>	Size of disk as presented to the VM, in bytes. Depending on the storage back-end type, the size may not be respected exactly	Read only
<code>physical-utilisation</code>	Amount of physical space that the VDI is taking up on the SR, in bytes	Read only
<code>type</code>	Type of VDI, for example, System or User	Read only
<code>sharable</code>	Value is true if this VDI may be shared, false otherwise	Read only
<code>read-only</code>	Value is true if this VDI can only be mounted read-only, false otherwise	Read only
<code>storage-lock</code>	Value is true if this VDI is locked at the storage level, false otherwise	Read only
<code>parent</code>	References the parent VDI when this VDI is part of a chain	Read only
<code>missing</code>	Value is true if SR scan operation reported this VDI as not present	Read only
<code>other-config</code>	Extra configuration information for this VDI	Read/write map parameter
<code>sr-name-label</code>	Name of the containing storage repository	Read only
<code>location</code>	Location information	Read only
<code>managed</code>	Value is true if the VDI is managed	Read only
<code>xenstore-data</code>	Data to be inserted into the <code>xenstore</code> tree (<code>/local/domain/0/backend/vbd/<i>domid</i>/<i>device-id</i>/smdata</code>) after the VDI is attached. The SM back-ends usually set this field on <code>vdi_attach</code> .	Read only map parameter
<code>sm-config</code>	SM dependent data	Read only map parameter

Parameter Name	Description	Type
<code>is-a-snapshot</code>	Value is true if this VDI is a VM storage snapshot, false otherwise	Read only
<code>snapshot_of</code>	The UUID of the storage this VDI is a snapshot of	Read only
<code>snapshots</code>	The UUIDs of all snapshots of this VDI	Read only
<code>snapshot_time</code>	The timestamp of the snapshot operation that created this VDI	Read only
<code>metadata-of-pool</code>	The uuid of the pool which created this metadata VDI	Read only
<code>metadata-latest</code>	Flag indicating whether the VDI contains the latest known metadata for this pool	Read only
<code>cbt-enabled</code>	Flag indicating whether changed block tracking is enabled for the VDI	Read/write

vdi-clone

```
1 xe vdi-clone uuid=uuid_of_the_vdi [driver-params:key=value]
```

Create a new, writable copy of the specified VDI that can be used directly. It is a variant of `vdi-copy` that is can expose high-speed image clone facilities where they exist.

Use the optional `driver-params` map parameter to pass extra vendor-specific configuration information to the back-end storage driver that the VDI is based on. For more information, see the storage vendor driver documentation.

vdi-copy

```
1 xe vdi-copy uuid=uuid_of_the_vdi sr-uuid=uuid_of_the_destination_sr
```

Copy a VDI to a specified SR.

vdi-create

```
1 xe vdi-create sr-uuid=uuid_of_sr_to_create_vdi_on name=label=  
   name_for_the_vdi type=system|user|suspend|crashdump virtual-size=  
   size_of_virtual_disk sm-config-\*=  
   storage_specific_configuration_data
```

Create a VDI.

The `virtual-size` parameter can be specified in bytes or using the IEC standard suffixes KiB, MiB, GiB, and TiB.

Note:

SR types that support thin provisioning of disks (such as Local VHD and NFS) do not enforce virtual allocation of disks. Take great care when over-allocating virtual disk space on an SR. If an over-allocated SR becomes full, disk space must be made available either on the SR target substrate or by deleting unused VDIs in the SR.

Some SR types might round up the `virtual-size` value to make it divisible by a configured block size.

vdi-data-destroy

```
1 xe vdi-data-destroy uuid=uuid_of_vdi
```

Destroy the data associated with the specified VDI, but keep the changed block tracking metadata.

Note:

If you use changed block tracking to take incremental backups of the VDI, ensure that you use the `vdi-data-destroy` command to delete snapshots but keep the metadata. Do not use `vdi-destroy` on snapshots of VDIs that have changed block tracking enabled.

vdi-destroy

```
1 xe vdi-destroy uuid=uuid_of_vdi
```

Destroy the specified VDI.

Note:

If you use changed block tracking to take incremental backups of the VDI, ensure that you use the `vdi-data-destroy` command to delete snapshots but keep the metadata. Do not use `vdi-destroy` on snapshots of VDIs that have changed block tracking enabled.

For Local VHD and NFS SR types, disk space is not immediately released on `vdi-destroy`, but periodically during a storage repository scan operation. If you must force deleted disk space to

be made available, call `sr-scan` manually.

vdi-disable-cbt

```
1 xe vdi-disable-cbt uuid=uuid_of_vdi
```

Disable changed block tracking for the VDI.

vdi-enable-cbt

```
1 xe vdi-enable-cbt uuid=uuid_of_vdi
```

Enable changed block tracking for the VDI.

Note:

You can enable changed block tracking only on licensed instances of XenServer Premium Edition.

vdi-export

```
1 xe vdi-export uuid=uuid_of_vdi filename=filename_to_export_to [format=
  format] [base=uuid_of_base_vdi] [--progress]
```

Export a VDI to the specified file name. You can export a VDI in one of the following formats:

- `raw`
- `vhd`

The VHD format can be *sparse*. If there are unallocated blocks within the VDI, these blocks might be omitted from the VHD file, therefore making the VHD file smaller. You can export to VHD format from all supported VHD-based storage types (EXT3/EXT4, NFS).

If you specify the `base` parameter, this command exports only those blocks that have changed between the exported VDI and the base VDI.

vdi-forget

```
1 xe vdi-forget uuid=uuid_of_vdi
```

Unconditionally removes a VDI record from the database without touching the storage back-end. In normal operation, use `vdi-destroy` instead.

vdi-import

```
1 xe vdi-import uuid=uuid_of_vdi filename=filename_to_import_from [format=  
=format] [--progress]
```

Import a VDI. You can import a VDI from one of the following formats:

- raw
- vhd

vdi-introduce

```
1 xe vdi-introduce uuid=uuid_of_vdi sr-uuid=uuid_of_sr name-label=  
name_of_new_vdi type=system|user|suspend|crashdump location=  
device_location_(varies_by_storage_type) [name-description=  
description_of_vdi] [sharable=yes|no] [read-only=yes|no] [other-  
config=map_to_store_misc_user_specific_data] [xenstore-data=  
map_to_of_additional_xenstore_keys] [sm-config=  
storage_specific_configuration_data]
```

Create a VDI object representing an existing storage device, without actually modifying or creating any storage. This command is primarily used internally to introduce hot-plugged storage devices automatically.

vdi-list-changed-blocks

```
1 xe vdi-list-changed-blocks vdi-from-uuid=first-vdi-uuid vdi-to-uuid=  
second-vdi-uuid
```

Compare two VDIs and return the list of blocks that have changed between the two as a base64-encoded string. This command works only for VDIs that have changed block tracking enabled.

For more information, see [Changed block tracking](#).

vdi-pool-migrate

```
1 xe vdi-pool-migrate uuid=VDI_uuid sr-uuid=destination-sr-uuid
```

Migrate a VDI to a specified SR, while the VDI is attached to a running guest. (Storage live migration)

For more information, see [Migrate VMs](#).

vdi-resize

```
1 xe vdi-resize uuid=vdi_uuid disk-size=new_size_for_disk
```

Change the size of the VDI specified by UUID.

vdi-snapshot

```
1 xe vdi-snapshot uuid=uuid_of_the_vdi [driver-params=params]
```

Produces a read-write version of a VDI that can be used as a reference for backup or template creation purposes or both. Use the snapshot to perform a backup rather than installing and running backup software inside the VM. The VM continues running while external backup software streams the contents of the snapshot to the backup media. Similarly, a snapshot can be used as a “gold image” on which to base a template. A template can be made using any VDIs.

Use the optional `driver-params` map parameter to pass extra vendor-specific configuration information to the back-end storage driver that the VDI is based on. For more information, see the storage vendor driver documentation.

A clone of a snapshot always produces a writable VDI.

vdi-unlock

```
1 xe vdi-unlock uuid=uuid_of_vdi_to_unlock [force=true]
```

Attempts to unlock the specified VDIs. If `force=true` is passed to the command, it forces the unlocking operation.

vdi-update

```
1 xe vdi-update uuid=uuid
```

Refresh the fields of the VDI object in the database.

VIF commands

Commands for working with VIFs (Virtual network interfaces).

The VIF objects can be listed with the standard object listing command (`xe vif-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

VIF parameters

VIFs have the following parameters:

- `uuid` (read only) the unique identifier/object reference for the VIF
- `vm-uuid` (read only) the unique identifier/object reference for the VM that this VIF resides on
- `vm-name-label` (read only) the name of the VM that this VIF resides on
- `allowed-operations` (read only set parameter) a list of the operations allowed in this state
- `current-operations` (read only set parameter) a list of the operations that are currently in progress on this VIF
- `device` (read only) integer label of this VIF, indicating the order in which VIF back-ends were created
- `MAC` (read only) MAC address of VIF, as exposed to the VM
- `MTU` (read only) Maximum Transmission Unit of the VIF in bytes.

This parameter is read-only, but you can override the MTU setting with the `mtu` key using the `other-config` map parameter. For example, to reset the MTU on a virtual NIC to use jumbo frames:

```
1 xe vif-param-set \  
2     uuid=<vif_uuid> \  
3     other-config:mtu=9000
```

- `currently-attached` (read only) true if the device is attached
- `qos_algorithm_type` (read/write) QoS algorithm to use
- `qos_algorithm_params` (read/write map parameter) parameters for the chosen QoS algorithm
- `qos_supported_algorithms` (read only set parameter) supported QoS algorithms for this VIF
- `MAC-autogenerated` (read only) Value is true if the MAC address of the VIF was automatically generated
- `other-config` (read/write map parameter) extra configuration `key:value` pairs
- `other-config:ethtoolrx` (read/write) set to on to enable receive checksum, off to disable
- `other-config:ethtooltx` (read/write) set to on to enable transmit checksum, off to disable
- `other-config:ethtoolsg` (read/write) set to on to enable scatter gather, off to disable

- `other-config:ethtooltso` (read/write) set to on to enable TCP segmentation offload, off to disable
- `other-config:ethtoolufo` (read/write) set to on to enable UDP fragment offload, off to disable
- `other-config:ethtoolgso` (read/write) set to on to enable generic segmentation offload, off to disable
- `other-config:promiscuous` (read/write) true to a VIF to be promiscuous on the bridge, so that it sees all traffic over the bridge. Useful for running an Intrusion Detection System (IDS) or similar in a VM.
- `network-uuid` (read only) the unique identifier/object reference of the virtual network to which this VIF is connected
- `network-name-label` (read only) the descriptive name of the virtual network to which this VIF is connected
- `io_read_kbs` (read only) average read rate in kB/s for this VIF
- `io_write_kbs` (read only) average write rate in kB/s for this VIF
- `locking_mode` (read/write) Affects the VIFs ability to filter traffic to/from a list of MAC and IP addresses. Requires extra parameters.
- `locking_mode:default` (read/write) Varies according to the default locking mode for the VIF network.

If the default-locking-mode is set to `disabled`, XenServer applies a filtering rule so that the VIF cannot send or receive traffic. If the default-lockingmode is set to `unlocked`, XenServer removes all the filtering rules associated with the VIF. For more information, see [Network Commands](#).

- `locking_mode:locked` (read/write) Only traffic sent to or sent from the specified MAC and IP addresses is allowed on the VIF. If no IP addresses are specified, no traffic is allowed.
- `locking_mode:unlocked` (read/write) No filters are applied to any traffic going to or from the VIF.
- `locking_mode:disabled` (read/write) XenServer applies a filtering rule is applied so that the VIF drops all traffic.

vif-create

```
1 xe vif-create vm-uuid=uuid_of_the_vm device=see below network-uuid=
   uuid_of_network_to_connect_to [mac=mac_address]
```

Create a VIF on a VM.

Appropriate values for the `device` field are listed in the parameter `allowed-VIF-devices` on the specified VM. Before any VIFs exist there, the values allowed are integers from 0–6.

The `mac` parameter is the standard MAC address in the form `aa:bb:cc:dd:ee:ff`. If you leave it unspecified, an appropriate random MAC address is created. You can also explicitly set a random MAC address by specifying `mac=random`.

vif-destroy

```
1 xe vif-destroy uuid=uuid_of_vif
```

Destroy a VIF.

vif-move

```
1 xe vif-move uuid=uuid network-uuid=network_uuid
```

Move the VIF to another network.

vif-plug

```
1 xe vif-plug uuid=uuid_of_vif
```

Attempt to attach the VIF while the VM is in the running state.

vif-unplug

```
1 xe vif-unplug uuid=uuid_of_vif
```

Attempts to detach the VIF from the VM while it is running.

vif-configure-ipv4

Configure IPv4 settings for this virtual interface. Set IPv4 settings as below:

```
1 xe vif-configure-ipv4 uuid=uuid_of_vif mode=static address=CIDR_address  
gateway=gateway_address
```

For example:


```
1 VIF.configure_ipv4(vifObject,"static", " 192.168.1.10/24", "
  192.168.1.1")
```

Clean IPv4 settings as below:

```
1 xe vif-configure-ipv4 uuid=uuid_of_vif mode=none
```

vif-configure-ipv6

Configure IPv6 settings for this virtual interface. Set IPv6 settings as below:

```
1 xe vif-configure-ipv6 uuid=uuid_of_vif mode=static address=IP_address
  gateway=gateway_address
```

For example:

```
1 VIF.configure_ipv6(vifObject,"static", "fd06:7768:b9e5:8b00::5001/64",
  "fd06:7768:b9e5:8b00::1")
```

Clean IPv6 settings as below:

```
1 xe vif-configure-ipv6 uuid=uuid_of_vif mode=none
```

VLAN commands

Commands for working with VLANs (virtual networks). To list and edit virtual interfaces, refer to the PIF commands, which have a VLAN parameter to signal that they have an associated virtual network. For more information, see [PIF commands](#). For example, to list VLANs, use `xe pif-list`.

vlan-create

```
1 xe vlan-create pif-uuid=uuid_of_pif vlan=vlan_number network-uuid=
  uuid_of_network
```

Create a VLAN on your XenServer host.

pool-vlan-create

```
1 xe pool-vlan-create pif-uuid=uuid_of_pif vlan=vlan_number network-uuid=
  uuid_of_network
```

Create a VLAN on all hosts in a pool by determining which interface (for example, `eno12429np0`) the specified network is on for each host, and then creating and plugging a new PIF object on each host accordingly.

vlan-destroy

```
1 xe vlan-destroy uuid=uuid_of_pif_mapped_to_vlan
```

Destroy a VLAN. Requires the UUID of the PIF that represents the VLAN.

VM commands

Commands for controlling VMs and their attributes.

VM selectors

Several of the commands listed here have a common mechanism for selecting one or more VMs on which to perform the operation. The simplest way is by supplying the argument `vm=name_or_uuid`. An easy way to get the uuid of an actual VM is to, for example, run `xe vm-list power-state=running`. (Get the full list of fields that can be matched by using the command `xe vm-list params=all`.) For example, specifying `power-state=halted` selects VMs whose `power-state` parameter is equal to `halted`. Where multiple VMs are matching, specify the option `--multiple` to perform the operation. The full list of parameters that can be matched is described at the beginning of this section.

The VM objects can be listed with the standard object listing command (`xe vm-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

VM parameters

VMs have the following parameters:

Note:

All writable VM parameter values can be changed while the VM is running, but new parameters are *not* applied dynamically and cannot be applied until the VM is rebooted.

- `appliance` (read/write) the appliance/vApp to which the VM belongs
- `uuid` (read only) the unique identifier/object reference for the VM
- `name-label` (read/write) the name of the VM
- `name-description` (read/write) the description string of the VM
- `order` (read/write) for vApp startup/shutdown and for startup after HA failover. VMs with an order value of 0 (zero) are started first, then VMs with an order value of 1, and so on.

- `version` (read only) the number of times this VM has been recovered. If you want to overwrite a new VM with an older version, call `vm-recover`
- `user-version` (read/write) string for creators of VMs and templates to put version information
- `is-a-template` (read/write) Value is false unless this VM is a template. Template VMs can never be started, they are used only for cloning other VMs After this value has been set to true it cannot be reset to false. Template VMs cannot be converted into VMs using this parameter.

You can convert a VM to a template with:

```
1 xe vm-param-set uuid=<vm uuid> is-a-template=true
```

- `is-control-domain` (read only) Value is true if this is a control domain (domain 0 or a driver domain)
- `power-state` (read only) current power state
- `start-delay` (read/write) the delay to wait before a call to start up the VM returns in seconds
- `shutdown-delay` (read/write) the delay to wait before a call to shut down the VM returns in seconds
- `memory-dynamic-max` (read/write) dynamic maximum in bytes
- `memory-dynamic-min` (read/write) dynamic minimum in bytes
- `memory-static-max` (read/write) statically set (absolute) maximum in bytes. If you want to change this value, the VM must be shut down.
- `memory-static-min` (read/write) statically set (absolute) minimum in bytes. If you want to change this value, the VM must be shut down.
- `suspend-VDI-uuid` (read only) the VDI that a suspend image is stored on
- `VCPUs-params` (read/write map parameter) configuration parameters for the selected vCPU policy.
- `VCPUs-max` (read/write) maximum number of virtual CPUs.
- `VCPUs-at-startup` (read/write) boot number of virtual CPUs
- `actions-after-crash` (read/write) action to take if the VM crashes. For PV guests, valid parameters are:
 - `preserve` (for analysis only)
 - `coredump_and_restart` (record a coredump and reboot VM)
 - `coredump_and_destroy` (record a coredump and leave VM halted)
 - `restart` (no coredump and restart VM)
 - `destroy` (no coredump and leave VM halted)

- `console-uuids` (read only set parameter) virtual console devices
- `platform` (read/write map parameter) platform-specific configuration

To disable VDA to switch Windows 10 into Tablet mode:

```
1 xe vm-param-set uuid=<vm_uuid> platform:acpi_laptop_slate=0
```

To enable VDA to switch Windows 10 into Tablet mode:

```
1 xe vm-param-set uuid=<vm_uuid> platform:acpi_laptop_slate=1
```

To check current state:

```
1 xe vm-param-get uuid=<vm_uuid> param-name=platform param-key=acpi_laptop_slate
```

- `allowed-operations` (read only set parameter) list of the operations allowed in this state
- `current-operations` (read only set parameter) a list of the operations that are currently in progress on the VM
- `allowed-VBD-devices` (read only set parameter) list of VBD identifiers available for use, represented by integers of the range 0–254. This list is informational only, and other devices may be used (but might not work).
- `allowed-VIF-devices` (read only set parameter) list of VIF identifiers available for use, represented by integers of the range 0–6. This list is informational only, and other devices may be used (but might not work).
- `HVM-boot-policy` (read/write) the boot policy for guests. Either BIOS Order or an empty string.
- `HVM-boot-params` (read/write map parameter) the order key controls the guest boot order, represented as a string where each character is a boot method: d for the CD/DVD, c for the root disk, and n for network PXE boot. The default is dc.
- `HVM-shadow-multiplier` (read/write) Floating point value which controls the amount of shadow memory overhead to grant the VM. Defaults to 1.0 (the minimum value), and only change this value if you are an advanced user.
- `PV-kernel` (read/write) path to the kernel
- `PV-ramdisk` (read/write) path to the `initrd`
- `PV-args` (read/write) string of kernel command line arguments
- `PV-legacy-args` (read/write) string of arguments to make legacy VMs boot
- `PV-bootloader` (read/write) name of or path to bootloader
- `PV-bootloader-args` (read/write) string of miscellaneous arguments for the bootloader

- `last-boot-CPU-flags` (read only) describes the CPU flags on which the VM was last booted
- `resident-on` (read only) the XenServer host on which a VM is resident
- `affinity` (read/write) The XenServer host which the VM has preference for running on. Used by the `xe vm-start` command to decide where to run the VM. Also known as the home server.
- `groups` (read/write) the VM group this VM belongs to. A VM group defines where XenServer attempts to start the VM or evacuate it to.
- `other-config` (read/write map parameter) A list of key-value pairs that specify extra configuration parameters for the VM.

For example, the `other-config` key-value pair `auto_poweron: true` requests to start the VM automatically after any host in the pool boots. You must also set this parameter in your pool's `other-config`. These parameters are now deprecated. Use the `ha-restart-priority` parameter instead.

- `start-time` (read only) timestamp of the date and time that the metrics for the VM were read. This timestamp is in the form `yyyymmddThh:mm:ss z`, where `z` is the single letter military timezone indicator, for example, `Z` for UTC (GMT)
- `install-time` (read only) timestamp of the date and time that the metrics for the VM were read. This timestamp is in the form `yyyymmddThh:mm:ss z`, where `z` is the single letter military timezone indicator, for example, `Z` for UTC (GMT)
- `memory-actual` (read only) the actual memory being used by a VM
- `VCPUs-number` (read only) the number of virtual CPUs assigned to the VM for a Linux VM. This number can differ from `VCPUs-max` and can be changed without rebooting the VM using the `vm-vcpu-hotplug` command. For more information, see `vm-vcpu-hotplug`. Windows VMs always run with the number of vCPUs set to `VCPUsmax` and must be rebooted to change this value. Performance drops sharply when you set `VCPUs-number` to a value greater than the number of physical CPUs on the XenServer host.
- `VCPUs-Utilization` (read only map parameter) a list of virtual CPUs and their weight
- `os-version` (read only map parameter) the version of the operating system for the VM
- `PV-drivers-version` (read only map parameter) the versions of the paravirtualized drivers for the VM
- `PV-drivers-detected` (read only) flag for latest version of the paravirtualized drivers for the VM
- `memory` (read only map parameter) memory metrics reported by the agent on the VM
- `disks` (read only map parameter) disk metrics reported by the agent on the VM
- `networks` (read only map parameter) network metrics reported by the agent on the VM

- `other` (read only map parameter) other metrics reported by the agent on the VM
- `guest-metrics-lastupdated` (read only) timestamp when the in-guest agent performed the last write to these fields. The timestamp is in the form `yyyymmddThh:mm:ss z`, where `z` is the single letter military timezone indicator, for example, `Z` for UTC (GMT)
- `actions-after-shutdown` (read/write) action to take after the VM has shutdown
- `actions-after-reboot` (read/write) action to take after the VM has rebooted
- `possible-hosts` potential hosts of this VM read only
- `dom-id` (read only) domain ID (if available, -1 otherwise)
- `recommendations` (read only) XML specification of recommended values and ranges for properties of this VM
- `xenstore-data` (read/write map parameter) data to be inserted into the `xenstore` tree (`/local/domain/*domid*/vm-data`) after the VM is created
- `is-a-snapshot` (read only) Value is true if this VM is a snapshot
- `snapshot_of` (read only) the UUID of the VM that this snapshot is of
- `snapshots` (read only) the UUIDs of all snapshots of this VM
- `snapshot_time` (read only) the timestamp of the snapshot operation that created this VM snapshot
- `memory-target` (read only) the target amount of memory set for this VM
- `blocked-operations` (read/write map parameter) lists the operations that cannot be performed on this VM
- `last-boot-record` (read only) record of the last boot parameters for this template, in XML format
- `ha-always-run` (read/write) Value is true if this VM is always restarted on another host if there is a failure of the host it is resident on. This parameter is now deprecated. Use the `ha-restart-priority` parameter instead.
- `ha-restart-priority` (read/write) restart or best-effort
- `blobs` (read only) binary data store
- `live` (read only) Value is true if the VM is running. Value is false if HA suspects that the VM is not be running.

vm-assert-can-be-recovered

```
1 xe vm-assert-can-be-recovered uuid [database] vdi-uuid
```

Tests whether storage is available to recover this VM.

vm-call-plugin

```
1 xe vm-call-plugin vm-uuid=vm_uuid plugin=plugin fn=function [args:key=value]
```

Calls the function within the plug-in on the given VM with optional arguments (args:key=value). To pass a "value" string with special characters in it (for example new line), an alternative syntax args:key:file=local_file can be used in place, where the content of local_file will be retrieved and assigned to "key" as a whole.

vm-cd-add

```
1 xe vm-cd-add cd-name=name_of_new_cd device=integer_value_of_an_available_vbd [vm-selector=vm_selector_value...]
```

Add a new virtual CD to the selected VM. Select the **device** parameter from 0–3.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-cd-eject

```
1 xe vm-cd-eject [vm-selector=vm_selector_value...]
```

Eject a CD from the virtual CD drive. This command only works if exactly one CD is attached to the VM. When there are two or more CDs, use the command **xe vbd-eject** and specify the UUID of the VBD.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-cd-insert

```
1 xe vm-cd-insert cd-name=name_of_cd [vm-selector=vm_selector_value...]
```

Insert a CD into the virtual CD drive. This command only works if there is exactly one empty CD device attached to the VM. When there are two or more empty CD devices, use the **xe vbd-insert** command and specify the UUIDs of the VBD and of the VDI to insert.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-cd-list

```
1 xe vm-cd-list [vbd-params] [vdi-params] [vm-selector=vm_selector_value ...]
```

Lists CDs attached to the specified VMs.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

You can also select which VBD and VDI parameters to list.

vm-cd-remove

```
1 xe vm-cd-remove cd-name=name_of_cd [vm-selector=vm_selector_value...]
```

Remove a virtual CD from the specified VMs.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-checkpoint

```
1 xe vm-checkpoint new-name-label=name_label [new-name-description=description]
```

Checkpoint an existing VM, using storage-level fast disk snapshot operation where available.

vm-clone

```
1 xe vm-clone new-name-label=name_for_clone [new-name-description=description_for_clone] [vm-selector=vm_selector_value...]
```

Clone an existing VM, using storage-level fast disk clone operation where available. Specify the name and the optional description for the resulting cloned VM using the **new-name-label** and **new-name-description** arguments.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-compute-maximum-memory

```
1 xe vm-compute-maximum-memory total=
  amount_of_available_physical_ram_in_bytes [approximate=add overhead
  memory for additional vCPUS? true|false] [vm-selector=
  vm_selector_value...]
```

Calculate the maximum amount of static memory which can be allocated to an existing VM, using the total amount of physical RAM as an upper bound. The optional parameter [approximate](#) reserves sufficient extra memory in the calculation to account for adding extra vCPUs into the VM later.

For example:

```
1 xe vm-compute-maximum-memory vm=testvm total=`xe host-list params=
  memory-free --minimal`
```

This command uses the value of the [memory-free](#) parameter returned by the `xe host-list` command to set the maximum memory of the VM named `testvm`.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-compute-memory-overhead

```
1 xe vm-compute-memory-overhead
```

Computes the virtualization memory overhead of a VM.

vm-copy

```
1 xe vm-copy new-name-label=name_for_copy [new-name-description=
  description_for_copy] [sr-uuid=uuid_of_sr] [vm-selector=
  vm_selector_value...]
```

Copy an existing VM, but without using storage-level fast disk clone operation (even if this option is available). The disk images of the copied VM are guaranteed to be *full images*, that is, not part of a copy-on-write (CoW) chain.

Specify the name and the optional description for the resulting copied VM using the [new-name-label](#) and [new-name-description](#) arguments.

Specify the destination SR for the resulting copied VM using the `sr-uuid`. If this parameter is not specified, the destination is the same SR that the original VM is in.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-copy-bios-strings

```
1 xe vm-copy-bios-strings host-uuid=host_uuid
```

Copy the BIOS strings of the given host to the VM.

Note:

After you first start a VM, you cannot change its BIOS strings. Ensure that the BIOS strings are correct before starting the VM for the first time.

vm-crashdump-list

```
1 xe vm-crashdump-list [vm-selector=vm selector value...]
```

List crashdumps associated with the specified VMs.

When you use the optional argument `params`, the value of `params` is a string containing a list of parameters of this object that you want to display. Alternatively, you can use the keyword `all` to show all parameters. If `params` is not used, the returned list shows a default subset of all available parameters.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-data-source-list

```
1 xe vm-data-source-list [vm-selector=vm selector value...]
```

List the data sources that can be recorded for a VM.

Select the VMs on which to perform this operation by using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all VMs.

Data sources have two parameters –`standard` and `enabled` –which you can see in the output of this command. If a data source has `enabled` set to **true**, the metrics are currently being recorded to the performance database. If a data source has `standard` set to **true**, the metrics are recorded to the performance database by default (and `enabled` is also set to **true** for this data source). If a data source has `standard` set to **false**, the metrics are *not* recorded to the performance database by default (and `enabled` is also set to **false** for this data source).

To start recording data source metrics to the performance database, run the `vm-data-source-record` command. This command sets `enabled` to **true**. To stop, run the `vm-data-source-forget`. This command sets `enabled` to **false**.

vm-data-source-record

```
1 xe vm-data-source-record data-source=name_description_of_data-source [
    vm-selector=vm selector value...]
```

Record the specified data source for a VM.

This operation writes the information from the data source to the persistent performance metrics database of the specified VMs. For performance reasons, this database is distinct from the normal agent database.

Select the VMs on which to perform this operation by using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all VMs.

vm-data-source-forget

```
1 xe vm-data-source-forget data-source=name_description_of_data-source [
    vm-selector=vm selector value...]
```

Stop recording the specified data source for a VM and forget all of the recorded data.

Select the VMs on which to perform this operation by using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all VMs.

vm-data-source-query

```
1 xe vm-data-source-query data-source=name_description_of_data-source [vm
    -selector=vm_selector_value...]
```

Display the specified data source for a VM.

Select the VMs on which to perform this operation by using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation is performed on all VMs.

vm-destroy

```
1 xe vm-destroy uuid=uuid_of_vm
```

Destroy the specified VM. This leaves the storage associated with the VM intact. To delete storage as well, use `xe vm-uninstall`.

vm-disk-add

```
1 xe vm-disk-add disk-size=size_of_disk_to_add device=uuid_of_device [vm-selector=vm_selector_value...]
```

Add a disk to the specified VMs. Select the `device` parameter from the value of the `allowed-VBD-devices` parameter of the VMs.

The `disk-size` parameter can be specified in bytes or using the IEC standard suffixes KiB, MiB, GiB, and TiB.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-disk-list

```
1 xe vm-disk-list [vbd-params] [vdi-params] [vm-selector=vm_selector_value...]
```

Lists disks attached to the specified VMs. The `vbd-params` and `vdi-params` parameters control the fields of the respective objects to output. Give the parameters as a comma-separated list, or the special key `all` for the complete list.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-disk-remove

```
1 xe vm-disk-remove device=integer_label_of_disk [vm-selector=
  vm_selector_value...]
```

Remove a disk from the specified VMs and destroy it.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-export

```
1 xe vm-export filename=export_filename [metadata=true|false] [vm-
  selector=vm_selector_value...]
```

Export the specified VMs (including disk images) to a file on the local machine. Specify the file name to export the VM into using the `filename` parameter. By convention, the file name has a `.xva` extension.

If the `metadata` parameter is `true`, the disks are not exported. Only the VM metadata is written to the output file. Use this parameter when the underlying storage is transferred through other mechanisms, and permits the VM information to be recreated. For more information, see [vm-import](#).

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-import

```
1 xe vm-import filename=export_filename [metadata=true|false] [preserve=
  true|false] [sr-uuid=destination_sr_uuid]
```

Import a VM from a previously exported file. If `preserve` is set to `true`, the MAC address of the original VM is preserved. The `sr-uuid` determines the destination SR to import the VM into. If this parameter is not specified, the default SR is used.

If the `metadata` is `true`, you can import a previously exported set of metadata without their associated disk blocks. Metadata-only import fails if any VDIs cannot be found (named by SR and [VDI location](#)) unless the `--force` option is specified, in which case the import proceeds regardless. If disks can be mirrored or moved out-of-band, metadata import/export is a fast way of moving VMs between disjoint pools. For example, as part of a disaster recovery plan.

Note:

Multiple VM imports are performed faster in serial than in parallel.

vm-install

```
1 xe vm-install new-name-label=name [template-uuid=
    uuid_of_desired_template] [template=template_uuid_or_name] [sr-uuid=
    sr_uuid | sr-name-label=name_of_sr] [copy-bios-strings-from=host_uuid
    ]
```

Install or clone a VM from a template. Specify the template name using either the `template-uuid` or `template` argument. Specify an SR using either the `sr-uuid` or `sr-name-label` argument. Specify to install BIOS-locked media using the `copy-bios-strings-from` argument.

Note:

When installing from a template that has existing disks, by default, new disks are created in the same SR as these existing disks. Where the SR supports it, these disks are fast copies. If a different SR is specified on the command line, the new disks are created there. In this case, a fast copy is not possible and the disks are full copies.

When installing from a template that doesn't have existing disks, any new disks are created in the SR specified, or the pool default SR when an SR is not specified.

vm-is-bios-customized

```
1 xe vm-is-bios-customized
```

Indicates whether the BIOS strings of the VM have been customized.

vm-memory-dynamic-range-set

```
1 xe vm-memory-dynamic-range-set min=min max=max
```

Configure the dynamic memory range of a VM. The dynamic memory range defines soft lower and upper limits for a VM's memory. It's possible to change these fields when a VM is running or halted. The dynamic range must fit within the static range.

vm-memory-limits-set

```
1 xe vm-memory-limits-set static-min=static_min static-max=static_max
    dynamic-min=dynamic_min dynamic-max=dynamic_max
```

Configure the memory limits of a VM.

vm-memory-set

```
1 xe vm-memory-set memory=memory
```

Configure the memory allocation of a VM.

vm-memory-shadow-multiplier-set

```
1 xe vm-memory-shadow-multiplier-set [vm-selector=vm_selector_value...] [
    multiplier=float_memory_multiplier]
```

Set the shadow memory multiplier for the specified VM.

This is an advanced option which modifies the amount of *shadow memory* assigned to a hardware-assisted VM.

In some specialized application workloads, such as Citrix Virtual Apps™, extra shadow memory is required to achieve full performance.

This memory is considered to be an overhead. It is separated from the normal memory calculations for accounting memory to a VM. When this command is invoked, the amount of free host memory decreases according to the multiplier and the `HVM_shadow_multiplier` field is updated with the value that Xen has assigned to the VM. If there is not enough XenServer host memory free, an error is returned.

The VMs on which to perform this operation are selected using the standard selection mechanism. For more information, see [VM selectors](#).

vm-memory-static-range-set

```
1 xe vm-memory-static-range-set min=min max=max
```

Configure the static memory range of a VM. The static memory range defines hard lower and upper limits for a VM's memory. It's possible to change these fields only when a VM is halted. The static range must encompass the dynamic range.

vm-memory-target-set

```
1 xe vm-memory-target-set target=target
```

Set the memory target for a halted or running VM. The given value must be within the range defined by the VM's `memory_static_min` and `memory_static_max` values.

vm-migrate

```
1 xe vm-migrate [compress=true|false] [copy=true|false] [host-uuid=
  destination_host_uuid] [host=name_or_ uuid_of_destination_host] [
  force=true|false] [live=true|false] [vm-selector=vm_selector_value
  ...] [remote-master=destination_pool_master_uuid] [remote-username=
  destination_pool_username] [remote-password=
  destination_pool_password] [remote-network=
  destination_pool_network_uuid] [vif:source_vif_uuid=
  destination_network_uuid] [vdi:vdi_uuid=destination_sr_uuid]
```

This command migrates the specified VMs between physical hosts.

The `compress` parameter overrides the `migration-compression` pool parameter for `xe pool-param-set`.

The `host` parameter in the `vm-migrate` command can be either the name or the UUID of the XenServer host. For example, to migrate the VM to another host in the pool, where the VM disks are on storage shared by both hosts:

```
1 xe vm-migrate uuid=vm_uuid host-uuid=destination_host_uuid
```

To move VMs between hosts in the same pool that do not share storage (storage live migration):

```
1 xe vm-migrate uuid=vm_uuid host-uuid=destination_host_uuid \
2   remote-master=192.0.2.35 remote-username=username remote-password=
   password
```

For storage live migration, you must provide the host name or IP address, user name, and password for the pool coordinator, even when you are migrating within the same pool.

You can choose the SR where each VDI gets stored:

```
1 xe vm-migrate uuid=vm_uuid remote-master=192.0.2.35 remote-username=
  username remote-password=password host-uuid=destination_host_uuid \
2   vdi:vdi_1=destination_sr1_uuid \
3   vdi:vdi_2=destination_sr2_uuid \
4   vdi:vdi_3=destination_sr3_uuid
```

Additionally, you can choose which network to attach the VM after migration:

```
1 xe vm-migrate uuid=vm_uuid \
2   vdi1:vdi_1_uuid=destination_sr1_uuid \
3   vdi2:vdi_2_uuid=destination_sr2_uuid \
4   vdi3:vdi_3_uuid=destination_sr3_uuid \
5   vif:source_vif_uuid=destination_network_uuid
```


For cross-pool migration:

```
1 xe vm-migrate uuid=vm_uuid remote-master=192.0.2.35 \  
2   remote-username=username remote-password=password \  
3   host-uuid=destination_host_uuid \  
4   vif:source_vif_uuid=destination_network_uuid \  
5   vdi:vdi_uuid=destination_sr_uuid
```

For more information about storage live migration, live migration, and live VDI migration, see [Migrate VMs](#).

Note:

If you are upgrading from an older version of XenServer, you might need to shut down and boot all VMs after migrating your VMs, to ensure that new virtualization features are picked up.

By default, the VM is suspended, migrated, and resumed on the other host. The [live](#) parameter selects live migration. Live migration keeps the VM running while performing the migration, thus minimizing VM downtime. In some circumstances, such as extremely memory-heavy workloads in the VM, live migration falls back into default mode and suspends the VM for a short time before completing the memory transfer.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-pause

```
1 xe vm-pause
```

Pause a running VM. Note this operation does not free the associated memory (see [vm-suspend](#)).

vm-reboot

```
1 xe vm-reboot [vm-selector=vm_selector_value...] [force=true]
```

Reboot the specified VMs.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Use the [force](#) argument to cause an ungraceful reboot. Where the shutdown is akin to pulling the plug on a physical server.

vm-recover

```
1 xe vm-recover vm-uuid [database] [vdi-uuid] [force]
```

Recovers a VM from the database contained in the supplied VDI.

vm-reset-powerstate

```
1 xe vm-reset-powerstate [vm-selector=vm_selector_value...] {  
2   force=true }
```

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

This is an *advanced* command only to be used when a member host in a pool goes down. You can use this command to force the pool coordinator to reset the power-state of the VMs to be **halted**. Essentially, this command forces the lock on the VM and its disks so it can be started next on another pool host. This call *requires* the force flag to be specified, and fails if it is not on the command-line.

vm-restart-device-models

```
1 xe vm-restart-device-models [vm-selector=vm_selector_value...]
```

Restart the device model for this VM on the host. While the device model is restarting, you can't stop, start, or migrate the VM. The end user of the VM might see a slight pause and resume in their session.

Note:

For the restart device model action to be supported on a Windows VM, the VM must have the XenServer VM Tools for Windows installed.

vm-resume

```
1 xe vm-resume [vm-selector=vm_selector_value...] [force=true|false] [on=  
   host_uuid]
```

Resume the specified VMs.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

If the VM is on a shared SR in a pool of hosts, use the [on](#) argument to specify which pool member to start it on. By default the system determines an appropriate host, which might be any of the members of the pool.

vm-retrieve-wlb-recommendations

```
1 xe vm-retrieve-wlb-recommendations
```

Retrieve the workload balancing recommendations for the selected VM.

vm-shutdown

```
1 xe vm-shutdown [vm-selector=vm_selector_value...] [force=true|false]
```

Shut down the specified VM.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Use the [force](#) argument to cause an ungraceful shutdown, similar to pulling the plug on a physical server.

vm-snapshot

```
1 xe vm-snapshot new-name-label=name_label [new-name-description+
  name_description]
```

Snapshot an existing VM, using storage-level fast disk snapshot operation where available.

vm-start

```
1 xe vm-start [vm-selector=vm_selector_value...] [force=true|false] [on=
  host_uuid] [--multiple]
```

Start the specified VMs.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

If the VMs are on a shared SR in a pool of hosts, use the [on](#) argument to specify which pool member to start the VMs on. By default the system determines an appropriate host, which might be any of the members of the pool.

vm-suspend

```
1 xe vm-suspend [vm-selector=vm_selector_value...]
```

Suspend the specified VM.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-uninstall

```
1 xe vm-uninstall [vm-selector=vm_selector_value...] [force=true|false]
```

Uninstall a VM, destroying its disks (those VDIs that are marked RW and connected to this VM only) in addition to its metadata record. To destroy just the VM metadata, use `xe vm-destroy`.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-unpause

```
1 xe vm-unpause
```

Unpause a paused VM.

vm-vcpu-hotplug

```
1 xe vm-vcpu-hotplug new-vcpus=new_total_vcpu_count [vm-selector=
vm_selector_value...]
```

Dynamically adjust the number of vCPUs available to a running Linux VM. The number of vCPUs is bounded by the parameter `VCPUs-max`. Windows VMs always run with the number of vCPUs set to `VCPUs-max` and must be rebooted to change this value.

Use the `new-vcpus` parameter to define the new *total* number of vCPUs that you want to have after running this command. Do not use this parameter to pass the number of vCPUs you want to add. For example, if you have two existing vCPUs in your VM and want to add two more vCPUs, specify `new-vcpus=4`.

The Linux VM or Windows VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Note:

When running Linux VMs without XenServer VM Tools installed, run the following command on the VM as `root` to ensure the newly hot plugged vCPUs are used: `# for i in /sys/devices/system/cpu/cpu[1-9]*/online; do if [ "$(cat $i)" = 0 ]; then echo 1 > $i; fi; done`

vm-vif-list

```
1 xe vm-vif-list [vm-selector=vm_selector_value...]
```

Lists the VIFs from the specified VMs.

The VM or VMs on which this operation is performed are selected using the standard selection mechanism. For more information, see [VM selectors](#). The selectors operate on the VM records when filtering, and *not* on the VIF values. Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Scheduled snapshots

Commands for controlling VM scheduled snapshots and their attributes.

The `vmss` objects can be listed with the standard object listing command (`xe vmss-list`), and the parameters manipulated with the standard parameter commands. For more information, see [Low-level parameter commands](#)

vmss-create

```
1 xe vmss-create enabled=true/false name=label=name type=type frequency=
  frequency retained-snapshots=value name-description=description
  schedule:schedule
```

Creates a snapshot schedule in the pool.

For example:

```
1 xe vmss-create retained-snapshots=9 enabled=true frequency=daily \
2   name-description=sample name=label=samplepolicy type=snapshot \
3   schedule:hour=10 schedule:min=30
```

Snapshot schedules have the following parameters:

Parameter Name	Description	Type
<code>name-label</code>	Name of the snapshot schedule.	Read/write
<code>name-description</code>	Description of the snapshot schedule.	Read/write
<code>type</code>	Disk snapshot or memory snapshot.	Read/write
<code>frequency</code>	Hourly; Daily; Weekly	Read/write
<code>retained-snapshots</code>	Snapshots to be retained. Range: 1-10.	Read/write
<code>schedule</code>	<code>schedule:days</code> (Monday to Sunday), <code>schedule:hours</code> (0 to 23), <code>schedule:minutes</code> (0, 15, 30, 45)	Read/write

vmss-destroy

```
1 xe vmss-destroy uuid=uuid
```

Destroys a snapshot schedule in the pool.

USB pass-through

USB pass-through is supported for the following USB versions: 1.1, 2.0, and 3.0.

USB pass-through enable/disable

```
1 xe pubsub-param-set uuid=pubsub_uuid passthrough-enabled=true/false
```

Enable/disable USB Pass-through.

pubsub-scan

```
1 xe pubsub-scan host-uuid=host_uuid
```

Scan PUSB and update.

vusb-create

```
1 xe vusb-create usb-group-uuid=usb_group_uuid vm-uuid=vm_uuid
```

Creates a virtual USB in the pool. Start the VM to pass through the USB to the VM.

vusb-unplug

```
1 xe vusb-unplug uuid=vusb_uuid
```

Unplugs USB from VM.

vusb-destroy

```
1 xe vusb-destroy uuid=vusb_uuid
```

Removes the virtual USB list from VM.

VM Group commands

Commands for controlling VM groups and their attributes.

VM Group parameters

- **uuid** (read only) the unique identifier/object reference for the VM
- **name-label** (read/write) the name of the VM
- **name-description** (read/write) the description string of the VM
- **placement** (read only) the type of placement rules to apply to VMs in the group. The available values are **anti-affinity** and **normal**.

vm-group-list

```
1 xe vm-group-list
```

Lists all VM groups in this pool.

vm-group-create

```
1 xe vm-group-create placement=anti-affinity|normal name-label=<name-label> [name-description=<name-description>]
```

Creates a VM group.

vm-group-destroy

```
1 xe vm-group-destroy uuid=<vm-group-uuid>
```

Deletes this VM group. If you delete a VM group that has VMs in it, the VMs are no longer associated with a group, but are not otherwise affected.

Infrastructure automation

December 17, 2025

Infrastructure automation enables you to manage and provision your IT resources more efficiently. By automating repetitive tasks, you reduce errors, increase consistency, and accelerate deployment processes. Unlike SDKs, which require coding expertise, tools like Terraform allow you to define and manage infrastructure declaratively, making management both accessible and reliable without needing deep programming skills.

Although this documentation currently focuses on the Terraform provider, XenServer® plans to expand its automation support to include other tools in the future.

Key benefits of infrastructure automation:

- **Consistency:** Automation ensures uniform configurations across environments, minimizing the risk of errors.
- **Efficiency:** Deploying new environments and configurations is faster and more streamlined, reducing the manual workload for IT teams.
- **Scalability:** Automation makes it easier to scale infrastructure reliably, essential for large-scale environments and rapidly growing organizations.

Several tools support infrastructure automation, each with unique capabilities. Currently, XenServer offers the Terraform Provider, which helps automate XenServer configuration and management tasks. For details, see [Automating XenServer with Terraform](#).

Automating XenServer® with Terraform

February 18, 2026

The XenServer Terraform Provider enables IT admins to automate configuration and management tasks for XenServer environments. This is especially useful in large-scale setups, where automation ensures consistency, reduces manual errors, and saves time. With the Terraform provider, you can:

- Create virtual machines (VMs)
- Manage snapshots (create and remove)
- Configure storage (add or remove local and NFS types)
- Set up and manage networks
- Retrieve information about XenServer hosts

Key benefits of the XenServer Terraform Provider:

- **Automation:** Simplifies the deployment and management of XenServer environments through Infrastructure-as-Code (IaC) practices.
- **Efficiency:** Automates tasks such as VM creation, snapshot management, and storage or network configuration, minimizing manual intervention.
- **Consistency:** Ensures uniform configurations across large environments, minimizing errors.
- **Scalability:** Supports rapid infrastructure growth by streamlining the creation and rebuilding of environments, such as QA setups.

Prerequisites

Before using the XenServer Terraform Provider, ensure that you have the following:

- Terraform version 1.8 or higher

Download Terraform from the official [Terraform Installation Guide](#).

Windows users can install Terraform using Chocolatey by running the following command:

```
1 choco install terraform
```

- XenServer instance with administrative access
- XenServer API credentials

These include the username, password, and the IP address of your XenServer.

Get started

1. Prepare the Terraform configuration.

To start, add the XenServer Terraform Provider to your Terraform configuration file `main.tf`:

```
1 terraform {  
2
```

```
3     required_providers {
4
5         xenserver = {
6
7             source = "xenserver/xenserver"
8         }
9     }
10 }
11
12 }
13
14
15 provider "xenserver" {
16
17     host      = "https://<ip address of XenServer>"
18     username  = "root"
19     password  = <root password>
20 }
```

2. Define resources.

For example, to create a VM, define the resource in your configuration file:

```
1 data "xenserver_sr" "sr" {
2
3     name_label = "Local storage"
4 }
5
6
7 resource "xenserver_vdi" "vdi1" {
8
9     name_label      = "local-storage-vdi-1"
10    sr_uuid          = data.xenserver_sr.sr.data_items[0].uuid
11    virtual_size     = 100 * 1024 * 1024 * 1024
12 }
13
14
15 data "xenserver_network" "network" {
16 }
17
18
19 resource "xenserver_vm" "windows_vm" {
20
21     name_label      = "Windows VM"
22     template_name   = "Windows 11"
23     static_mem_max  = 4 * 1024 * 1024 * 1024
24     vcpus           = 4
25     cores_per_socket = 2
26     cdrom           = "win11-x64_uefi.iso"
27     boot_mode       = "uefi_security"
28     boot_order      = "cdn"
29
30     hard_drive = [
31         {
```

```
32         vdi_uuid = xenserver_vdi.vdi1.uuid,
33         bootable = true,
34         mode      = "RW"
35     }
36 }
37
38 ]
39
40 network_interface = [
41     {
42         device      = "0"
43         network_uuid = data.xenserver_network.network.data_items
44             [0].uuid,
45     }
46 ]
47
48 other_config = {
49     "tf_created" = "true"
50 }
51
52 }
53
54 }
```

Or, to create an NFS Storage Repository (SR):

```
1 resource "xenserver_sr" "nfs" {
2
3     name_label    = "Test NFS SR"
4     type          = "nfs"
5     content_type  = ""
6     shared        = true
7     device_config = {
8
9         server      = "192.0.2.1"
10        serverpath   = "/server/path"
11        nfsversion   = "3"
12    }
13
14    sm_config = {
15
16        shared = "true"
17    }
18
19 }
```

3. Initialize and apply the configuration.

To install the XenServer Terraform Provider, run:

```
1 terraform init
```

Then, to provision the defined resources, run:

```
1 terraform apply
```

For advanced configurations and additional resources, refer to the [XenServer Terraform Provider Documentation](#). If you encounter any issues, contact Support or raise a GitHub issue ticket.

What's new

The XenServer Terraform Provider is updated independently of XenServer versions to introduce new automation capabilities, fixes, and ensure ongoing compatibility.

XenServer Terraform Provider 0.2.2

Released Mar 27, 2025

This release contains the following improvements:

- Update the Terraform provider dependencies to fix a security advisory in an upstream package.
- Release Terraformprovider with Go SDK v25.3.0.

Troubleshooting

February 18, 2026

If you experience technical difficulties with the XenServer® host, this section is meant to help you solve the problem if possible. If it isn't possible, use the information in this section to gather the application logs and other data that can help Technical Support track and resolve the issue.

The following articles provide troubleshooting information about specific areas of the product:

- [VM troubleshooting](#)
- [Networking troubleshooting](#)
- [Clustered pool troubleshooting](#)
- [XenCenter troubleshooting](#)
- [Workload Balancing troubleshooting](#)
- [Conversion Manager troubleshooting](#)

Troubleshoot connections between XenCenter® and the XenServer host

If you have trouble connecting to the XenServer host with XenCenter, check the following:

- Is your XenCenter an older version than the XenServer host that you are attempting to connect to?

XenCenter 8.2.7 and earlier are not supported with XenServer 9 hosts. To manage your XenServer 9 hosts or pools, you require the latest version of XenCenter with a version of the form YYYY.x.x.

To correct this issue, install the [latest version of XenCenter](#).

- Is your license current?

You can see the expiration date for your license access code in the XenServer host **General** tab under the **License Details** section in XenCenter.

For more information on licensing a host, see [Licensing](#).

- The XenServer host talks to XenCenter using HTTPS over the following ports:
 - Port 443 (a two-way connection for commands and responses using the management API)
 - Port 5900 for graphical VNC connections with paravirtualized Linux VMs.

If you have a firewall enabled between the XenServer host and the machine running the client software, ensure that it allows traffic from these ports. For more information, see [Internet connectivity](#).

Ping the XenServer host

If your XenServer host is not responsive to XenCenter, it might still be running but management functionality has failed. To check whether this is the case, complete the following steps:

1. Note the IP address or host name of the XenServer host.
2. Open a command line or terminal on a system that has network access to your host.
3. Run the following command: `ping <ip_address_or_host_name>`
 - If the command receives a response and SSH access is enabled on your host, you can SSH into your host and use the command line to gather diagnostic logs. For more information, see XenServer server status reports.
 - If the command does not receive a response, you might need to access the host through the serial console to gather diagnostic logs. For more information, see [Configure serial console access](#).

Gather logs

XenServer server status reports

Important:

XenServer server status reports can contain sensitive information. For more information, see [Data governance](#).

XenCenter can be used to gather XenServer host information:

1. In XenCenter, open the **Tools** menu and click **Server Status Report**. The **Server Status Report** task opens.
2. Select the servers to gather information from. Click **Next**.
3. Select from a list of different types of information (various logs, crash dumps, and so on) to include. Click **Next**.
4. Specify the file name and the location on the system running XenCenter to save the information to. Click **Next**.
5. Wait while the information is compiled and downloaded to the machine that XenCenter is running on.
6. Click **Finish** to exit the wizard.

You can also gather the server status report for a host by using the command line. If you cannot connect to your XenServer host by using XenCenter but can connect to it by SSH, you can use the server status report command line tool to gather the logs for that host.

To gather the log files, run the following command in the XenServer host console:

```
1 xenserver-status-report
```

By default, the files gathered for a server status report can be limited in size. If you need log files that are larger than the default, you can run the following command in the XenServer host console:

```
1 xenserver-status-report -u
```

After gathering the log files on your XenServer host, use a secure copy program, such as `pscp` or `WinSCP`, to copy the log files to your local system.

If your XenServer host is not booting successfully, you can still attempt to gather logs. For more information, see [Gather logs when XenServer is failing to boot](#).

Send syslog messages to a central server

In addition to having logs written to the control domain filesystem, you can configure your XenServer host to write them to a remote server. The remote server must have the `syslogd` daemon running

on it to receive the logs and aggregate them correctly. The `syslogd` daemon is a standard part of all flavors of Linux and Unix, and third-party versions are available for Windows and other operating systems.

Set the `syslog_destination` parameter to the host name or IP address of the remote server (and optionally specifying a port if you do not wish to use the default 514) where you want the logs to be written:

```
1 xe host-param-set uuid=host_uuid logging:syslog_destination=hostname[:port]
```

Run the command:

```
1 xe host-syslog-reconfigure host-uuid=host_uuid
```

To enforce the change. (You can also run this command remotely by specifying the `host` parameter.)

XenCenter logs

XenCenter also has a client-side log. This file includes a complete description of all operations and errors that occur when using XenCenter. It also contains informational logging of events that provide you with an audit trail of various actions that have occurred. The XenCenter log file is stored in your profile folder at the following path: `%userprofile%\AppData\Roaming\XenServer\XenCenter\logs\XenCenter.log`.

To locate the XenCenter log files - for example, when you want to open or email the log file - click **View XenCenter Log Files** in the XenCenter **Help** menu.

Installation logs

If you experience an unknown error during installation, capture the log file from your host and provide it to Technical Support.

Using a keyboard connected directly to the host machine (not connected over a serial port), you can access three virtual terminals during installation:

- Press **Alt+F1** to access the main XenServer Installer
- Press **Alt+F2** to access a local shell
- Press **Alt+F3** to access the event log

To capture and save the log files:

1. Press **Alt+F2** to access the local shell.

2. Enter the following:

```
1 /opt/xensource/installer/report.py
```

3. You are prompted to choose where you want to save the log file: **NFS**, **FTP**, or **Local media**.

Select **NFS** or **FTP** to copy the log file to another machine on your network. To do so, networking must be working properly, and you must have write access to a remote machine.

Select **Local media** to save the file to a removable storage device, such as a USB flash drive, on the local machine.

Once you have made your selections, the program writes the log file to your chosen location. The file name is `support.tar.bz2`.

Advanced troubleshooting

June 22, 2026

This article lists advanced troubleshooting actions that you can take with the guidance of the Technical Support team.

Change options on the Xen® command line

XenServer® is based on the open-source Xen hypervisor. As part of diagnosing or working around issues, you might be asked to add parameters to the Xen command line.

We strongly recommend that you only change your Xen command line under the guidance of Technical Support. Setting or changing any options that are not listed in the XenServer product documentation or given by Technical Support is not supported. In addition to rendering your environment ineligible for support, these changes to Xen command line arguments can make your hosts unbootable or potentially vulnerable to security threats.

To add options to the Xen command line of your host, complete the following steps:

1. Log in to the XenServer host console, either by using SSH or the **Console** tab in XenCenter.
2. Run the following command with the parameters and values advised by Technical Support: `/opt/xensource/libexec/xen-cmdline --set-xen <parameter>=<value>`
3. The command returns silently, but you can verify that your parameters are set by checking your `grub.cfg` file.
 - If you are using legacy BIOS boot mode for your host, run the following command: `more /boot/grub/grub.cfg`

- If you are using UEFI boot mode for your host, run the following command: `more /boot/efi/EFI/xenserver/grub.cfg`

The parameter you specified appears in the `XenServer` and `XenServer (Serial)` sections of the file.

4. Reboot your XenServer host to make your changes take effect.

To remove options from the Xen command line of your host, complete the following steps:

1. Log in to the XenServer host console, either by using SSH or the **Console** tab in XenCenter.
2. Run the following command with the parameters and values advised by Technical Support: `/opt/xensource/libexec/xen-cmdline --remove-xen <parameter>`
3. The command returns silently, but you can verify that your parameters are removed by checking your `grub.cfg` file.
 - If you are using legacy BIOS boot mode for your host, run the following command: `more /boot/grub/grub.cfg`
 - If you are using UEFI boot mode for your host, run the following command: `more /boot/efi/EFI/xenserver/grub.cfg`
4. Reboot your XenServer host to make your changes take effect.

Gather logs when XenServer is failing to boot

If your XenServer host is not booting successfully, you can still gather logs from it by attempting to boot from the installation ISO.

1. Download the installation ISO from <https://xenserver.com/downloads>.
2. Boot your XenServer host from the installation ISO. For information about methods of booting from the ISO, see step 2 in [Install the XenServer host](#).
3. When the first dialog appears on the screen, press **Alt+F2**. The installation process switches you to a shell console.
4. Verify whether the installation ISO can see the local or remote disk. Run the following command:

```
1 fdisk -l
```

If the local and remote disks are visible, this command lists the visible partitions. These partitions usually have names of the form `/dev/sdaX`.

5. Identify the root partition on the local disk. The root partition is usually `/dev/sda1`, but can be different on different hardware.

6. Mount the root partition in a temporary mount point `/mnt` to access the local disk:

```
1 mkdir /mnt
2 mount /dev/sda1 /mnt
3 ls /mnt
```

7. Change the temporary root file system to the file system on your disk:

```
1 chroot /mnt
```

8. Create a server status report by running the following commands:

```
1 mount -t sysfs none /sys
2 xenserver-status-report --yestoall
```

The command prints the location of the tarball containing the server status report.

9. Prepend `/mnt` to the path given by the previous step to get the path to the file in the temporary root file system, for example: `/mnt/var/opt/xen/bug-report/bug-report-xxxx.tar.bz2`
10. Set up a temporary IP address to copy the server status report over the network to your desktop.
- If you have DHCP running in the network, run the following command:
- ```
1 dhclient <network-interface>
```
- If there is no DHCP, manually configure the IP address:
- ```
1 ifconfig <network-interface> <x.x.x.x> netmask <x.x.x.x> up
2 route add default <x.x.x.x>
```
11. Using an SCP client, connect to the previously configured IP address and copy the status report to your local desktop.

Configure serial console access

Serial console access to the XenServer host can be useful for debug purposes. You can use the serial connection with HyperTerminal (or a similar application) to collect XenServer logs.

You can use a physical port to connect to your host. If you are using this method, you need the following things:

- 9-pin Serial Null Modem Cable.
- Serial COM Port on a system that is running software for connecting to the serial port, for example, Microsoft HyperTerminal.
- Serial COM Port on a XenServer host.

If your hosts do not have a physical serial port or suitable physical infrastructure is not available, investigate whether you can configure an embedded management device with your hardware. This method is also known as Serial over LAN. If you are using this method, you need the following things:

- A virtual serial port provided by a baseboard management controller such as Dell iDRAC or HP iLO.
- A computer with appropriate software to connect to the virtual serial port, for example, `ipmitool`.

Note:

By default, XenServer assumes a serial baud rate of 115200. If a different baud rate is required, adjust the value of `<speed>` in the `grub.cfg` file. If needed, seek the guidance of the Technical Support team.

To set up a serial connection, complete the following steps:

1. If you are not using COM1 on your server, edit your `grub.cfg` file.
 - For legacy BIOS boot the `grub.cfg` file is in the `/boot/grub` directory.
 - For UEFI boot the `grub.cfg` file is in the `/boot/efi/EFI/xenserver` directory.
2. Locate the following line: `serial --unit=0 --speed=<speed>`
3. Edit the value of `--unit` to be `n-1`, where `<n>` represents your COM port number.

For example, to use COM2, set the value of `--unit` to be 1:

```
1 serial --unit=1 --speed=<speed>
```

4. Locate the line in the `menuentry 'XenServer (Serial)'` section where the following text is displayed: `com1=<speed>,8n1 console=com1,vga`
5. Replace `com1` with `com<n>`, where `<n>` represents your COM port number.
6. Start or reboot the XenServer host.
7. From the grub menu, select **XenServer (Serial)**.
8. XenServer starts with serial console support.

You can now use your chosen method to connect to the serial console.

Generate a crash dump from the serial console

If a XenServer host crashes, the running kernel can migrate into a special memory area that is based on `kexec` functions. It loads a special kernel without the need for a server restart. This new kernel

attempts to gather as much detail about the crash as possible (memory dumps, process list) and save it into the `/var/crash` directory.

Sometimes, XenServer hangs and does not reboot after a crash. When this situation happens, the server is unresponsive but still running and only a reboot of the server can resolve the issue. In addition, there might be no crash dumps available on the server: there is no dump in the `/var/crash` directory and the `xe host-crashdump-list` contains no information about the crash.

In this situation, you can force a crashdump that saves information about the crash, which can then be analyzed.

After switching into the serial console, complete the following steps:

1. Wait until XenServer becomes unresponsive.
2. Press **Ctrl+A** three times to switch to the Xen hypervisor menu. You see this message on the serial console: `*** Serial input to Xen (type 'CTRL-a'three times to switch input)`
3. Press **h** to display all available operations.
4. Press **Shift+C** to trigger a crashdump.
5. Wait until the XenServer host reboots and save all data on a local file system.
6. Check for the new crash dump on the server in the `/var/crash` directory.

Switch to the debug hypervisor

In some situations, it might be useful to switch the host to using the debug hypervisor and attempt to reproduce the issue. The debug hypervisor performs more built-in checks and is more verbose in its output.

To enable the debug hypervisor on your host, complete the following steps:

1. Log in to the XenServer host console, either by using SSH or the **Console** tab in XenCenter.
2. Change directory to the `/boot` directory: `cd /boot`
3. Verify that you have the debug module `xen-debug.efi` file available: `ls -l xen*`

This command shows the available hypervisor versions and the symbolic links that are used when booting the host. In normal operation, `xen-release.efi` links to the release hypervisor file, `xen-<version>.efi`, and `xen-debug.efi` links to the debug hypervisor file, `xen-<version>-d.efi`. The debug hypervisor is the same version as the release hypervisor, but has an additional `-d` included in the file name.

```
xen-4.17.5-6.config
xen-4.17.5-6-d.config
xen-4.17.5-6-d.gz
xen-4.17.5-6-d.map
xen-4.17.5-6.gz
xen-4.17.5-6.map
xen-debug.gz -> xen-4.17.5-6-d.gz
xen-fallback.gz
xen.gz -> xen-4.17.5-6.gz
xen-release.gz -> xen-4.17.5-6.gz
```

4. Change the `xen.efi` link to point to the debug hypervisor: `ln -sf /boot/xen-debug.efi /boot/xen.efi`
5. Reboot your host.

After using the debug hypervisor to perform the actions recommended by Technical Support and gather more diagnostic information, switch your hypervisor back to the release hypervisor before returning to normal operation.

To return to the using the release hypervisor on your host, complete the following steps:

1. Log in to the XenServer host console, either by using SSH or the **Console** tab in XenCenter.
2. Change directory to the `/boot` directory: `cd /boot`
3. Change the `xen.efi` link to point to the release hypervisor: `ln -sf /boot/xen-release.efi /boot/xen.efi`
4. Reboot your host.

Support

July 2, 2026

We provide Technical Support services to customers with a XenServer license. To access this support, you can open a Support Case online or contact the support center by phone if you experience technical difficulties. For more information, see [the XenServer support page](#).

- If you are using a preview feature in XenServer®, you cannot access this support, but we do value your feedback. For more information, see [Support for preview features](#).
- If you are using XenServer Trial Edition (unlicensed), you cannot access this support, but we do value your feedback. For more information, see [Provide feedback for XenServer and XenCenter®](#).

Frequent updates

XenServer 9 uses a frequent update model that delivers features, fixes, and improvements to your hosts. We expect you to consume all outstanding updates within six months of their release to remain in support. If the update level of your pool is older than six months, we might ask you to reproduce the issue on the latest update level as part of our investigation.

Note:

Applying updates is only assured to work when the previous state of the pool is less than six months old. If you haven't applied updates for more than six months, updating to the latest update level is not tested and is not assured.

Bugs

Support checklist

This section guides you through possible actions to take when you experience an issue in your XenServer environment. By completing as many of these steps as you can, you help us troubleshoot your issue faster.

First steps

When first experiencing the issue, complete these steps:

1. Before carrying out any recovery steps, capture any logs that you can from the environment:
 - If the issue relates to XenServer, capture a server status report (SSR) of the host or pool where the issue was seen. Go to the XenCenter **Tools** menu then **Server Status Report**.
 - If the issue relates to XenCenter:
 - Get the application log files by going to the XenCenter **Help** menu then **View XenCenter Log Files**.
 - Capture screenshots of the relevant displays.
 - If the issue relates to a VM, gather any relevant logs from the VM operating system.

For more information about getting log information, see [Gather logs](#).

2. Make a note of the synchronization date, synchronization checksum, and update checksum for your host or pool. For more information, see [View pending tasks](#).
3. If necessary, attempt to recover your environment to a working state.

Self-help

We provide information and guidance that can help you diagnose and resolve issues you might experience.

1. Check this documentation for help:
 - [Known Issues](#): This article lists known issues in XenServer and, if applicable, any workarounds you can apply.
 - [Early Access](#) and [Normal](#): These articles list available updates to XenServer. A fix for your issue might have recently been released.
 - [Troubleshooting](#): This article is an entrypoint into the troubleshooting information provided in the documentation.
 - Review the section of the documentation related to the feature that is showing the issue. There might be constraints to this feature that are causing the issue or configuration options that can help fix it.
2. The [Citrix Knowledge Center](#) contains many articles written by our Technical Support team that describe solutions to previously seen issues with XenServer.

If you take any diagnostic action or change your environment configuration as a result of this information, make a note of it and let us know if you do raise an issue.

Update to the latest version

If you are already using the latest version of all your XenServer components, skip to Capture logs.

If you are not running the latest XenServer, XenCenter, or related component, the fix for your issue might be included in the latest version or update level. We recommend that you update your environment to the latest version, if possible.

1. [Update your pool to the latest update](#).
2. [Update XenCenter to the latest version](#).
3. If the issue relates to a Windows VM, [update the XenServer VM Tools for Windows to the latest version](#).
4. If the issue relates to a Linux VM, ensure the latest version of the XenServer VM Tools for Linux is installed. These tools are available on <https://xenserver.com/downloads>.
5. If the issue relates to Workload Balancing or the Conversion Manager, ensure you are using the latest version available on <https://xenserver.com/downloads>.

Reproduce the issue on the latest version

If you updated any of the components in your environment since you encountered the issue, attempt to reproduce the issue now.

Capture logs

- If the issue relates to XenServer, capture a server status report (SSR) of the host or pool where the issue was encountered. Go to the XenCenter **Tools** menu then **Server Status Report**.
- If the issue relates to XenCenter:
 - Get the application log files by going to the XenCenter **Help** menu then **View XenCenter Log Files**.
 - Capture screenshots of the relevant displays.
- If the issue relates to a VM, gather any relevant logs from the VM operating system.

Contact support

For methods of getting in touch with us, see [the XenServer support page](#).

Provide the following information to support:

- Date and time when the issue was encountered
- Pool update channel (Early Access or Normal)
- Pool coordinator's name
- Date and time when the pool was last synchronized
- Synchronization checksum for the pool coordinator
- Date and time when the pool was last updated
- Update checksum for all hosts in the pool
- Any relevant screenshots to help to describe the issue
- Any change or event that might have triggered the issue
- Any known workaround
- Any diagnostic steps that you've already taken
- The SSR for the problem pool
- If applicable, the XenCenter logs

Note:

If you are unable to reproduce the issue on an up-to-date pool, attach the SSR that you captured when the issue first happened and explain why you were unable to apply all updates.

Support for preview features**Note:**

Preview features can be changed or removed with minimal notice. Preview features are not guaranteed to be delivered to production support.

Public previews

Some updates include preview features, which give you an early look at capabilities that might be included in the product later. These preview features can be deployed in a production environment, but are not supported for production use.

The status of a preview feature can be indicated in the following ways:

- In the Updates View of XenCenter, the feature is listed as a Preview Feature.
- In product documentation for the feature, it is noted as (Preview).
- If the feature can be used through XenCenter, the UI text might note that it is a preview.

Some preview features might require you to set a flag in your environment to use the feature. For more information, see the product documentation for the feature.

If you have feedback about or encounter an issue with a preview feature, reach out to us at feedback@xenserver.com. Do not contact Technical Support for preview features.

Private previews

Sometimes, we provide preview features to a select set of customers as a private preview. If you are interested in working with us to try out features you might like to see in XenServer in future, get in touch at feedback@xenserver.com.

Private previews are not supported in your production environment.

Provide feedback for XenServer and XenCenter

Help us to improve our product by providing feedback on the features and usability of our new release. To provide feedback, do not contact Technical Support but instead submit a feedback email.

Submit a feedback email

Email any feedback and queries to feedback@xenserver.com. To help us understand the full context of your situation, ensure that you include the following information in your feedback email:

- Your full name
- Your company or business
- Your geographical location
- Your license type
- The number of hosts in your production deployment
- The guest OS where the issue occurred (Windows or Linux)

Do not use this email address to request Technical Support.

Third party notices

December 17, 2025

This release of XenServer® includes third-party software licensed under a number of different licenses.

To extract the licensing information from your installed XenServer product and components, see the instructions in [XenServer Open Source Licensing and Attribution](#).

In addition, note the following information:

- This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)
- This product includes cryptographic software written by Eric Young (eyay@cryptsoft.com).
- XenServer High Availability is powered by everRun, a registered trademark of Stratus Technologies Bermuda, Limited.

XenServer® Open Source Licensing and Attribution

December 17, 2025

The XenServer product is a compilation of software packages. Each package is governed by its own license. The complete licensing terms applicable to a given package can be found in the source RPM of the package, unless the package is covered by a proprietary license which does not permit source redistribution, in which case no source RPM is made available.

The XenServer distribution contains content from CentOS Linux and CentOS Stream. Where the CentOS Project holds any copyright in the packages making up the CentOS Linux or CentOS Stream distributions, that copyright is licensed under the GPLv2 license unless otherwise noted. For more information, see <https://www.centos.org/legal/licensing-policy/>.

Extracting attribution and licensing information on an installed XenServer host

This article provides a method to extract the licensing information from all RPM packages included in your XenServer installation.

Get overview information

To list all RPMs and their licenses:

1. Connect to your XenServer host console by SSH or through XenCenter.
2. At the console command line, run the following command:

```
1 rpm -qa --qf '%{
2   name }
3   -%{
4   version }
5   : %{
6   license }
7   \n'
```

This command lists all installed components and the licenses they are distributed under. The output is of the following form:

```
1 readline-6.2: GPLv3+
2 gnupg2-2.0.22: GPLv3+
3 libdb-5.3.21: BSD and LGPLv2 and Sleepycat
4 rpm-python-4.11.3: GPLv2+
5 sqlite-3.7.17: Public Domain
6 qrencode-libs-3.4.1: LGPLv2+
7 libselinux-2.5: Public Domain
8 ustr-1.0.4: MIT or LGPLv2+ or BSD
9 gdbm-1.10: GPLv3+
10 procps-ng-3.3.10: GPL+ and GPLv2 and GPLv2+ and GPLv3+ and LGPLv2+
11 p11-kit-trust-0.23.5: BSD
12 device-mapper-libs-1.02.149: LGPLv2
13 xenserver-release-8.2.50: GPLv2
14 elfutils-libs-0.170: GPLv2+ or LGPLv3+
15 xz-libs-5.2.2: LGPLv2+
16 dbus-1.10.24: (GPLv2+ or AFL) and GPLv2+
17 elfutils-libelf-0.170: GPLv2+ or LGPLv3+
18 systemd-sysv-219: LGPLv2+
19 jemalloc-3.6.0: BSD
```

Get detailed information

To obtain a more complete list of information about each installed component:

1. Connect to your XenServer host console by SSH or through XenCenter.
2. At the console command line, run the following command:

```
1 rpm -qai | sed '/^Name /i\\n'
```

The output is of the following form:

```
1 Name: host-upgrade-plugin
2 Version      : 2.2.6
3 Release      : 1.xs8
4 Architecture: noarch
5 Install Date: Wed 23 Aug 2023 01:54:25 PM UTC
6 Group: Unspecified
7 Size: 101626
8 License      : GPL
9 Signature    : RSA/SHA256, Tue 30 May 2023 10:01:44 AM UTC, Key ID
                5259d0b0f6529a4e
10 Source RPM  : host-upgrade-plugin-2.2.6-1.xs8.src.rpm
11 Build Date  : Fri 26 May 2023 03:05:49 AM UTC
12 Build Host  : cf27e1dd25c54cbb8cef79726ed2bf2c
13 Relocations : (not relocatable)
14 Packager    : Koji
15 Vendor      : Cloud Software Group, Inc.
16 Summary     : Host upgrade plugin
17 Description :
18 Host upgrade plugin.
19
20 Name        : m4
21 Version     : 1.4.16
22 Release     : 10.el7
23 Architecture: x86_64
24 Install Date: Wed 23 Aug 2023 01:52:31 PM UTC
25 Group       : Applications/Text
26 Size        : 525707
27 License     : GPLv3+
28 Signature   : RSA/SHA256, Tue 09 May 2023 02:53:25 PM UTC, Key ID
                5259d0b0f6529a4e
29 Source RPM  : m4-1.4.16-10.el7.src.rpm
30 Build Date  : Fri 20 Nov 2015 07:28:07 AM UTC
31 Build Host  : worker1.bsys.centos.org
32 Relocations : (not relocatable)
33 Packager    : CentOS BuildSystem <http://bugs.centos.org>
34 Vendor      : CentOS
35 URL         : http://www.gnu.org/software/m4/
36 Summary     : The GNU macro processor
37 Description :
38 A GNU implementation of the traditional UNIX macro processor. M4
    is
```

```
39 useful for writing text files which can be logically parsed, and
    is used
40 by many programs as part of their build process. M4 has built-in
41 functions for including files, running shell commands, doing
    arithmetic,
42 etc. The autoconf program needs m4 for generating configure
    scripts, but
43 not for running configure scripts.
```

Multiple licenses Some components in the XenServer product contain multiple licenses. For example, `procps-ng-3.3.10` contains the following parts:

- some parts which are licensed with the original GPL (or any later version)
- some parts which are licensed with the GPL version 2 (only)
- some parts which are licensed with the GPL version 2 (or any later version)
- some parts which are licensed with the GPL version 3 (or any later version)
- some parts which are licensed with the LGPL version 2 (or any later version)

In this case, inspect the documentation in `/usr/share/doc/procps-ng-3.3.10` for further information or, if necessary, the corresponding source RPM.

Get more information

In most cases, further information about each component and full license text is installed in either `/usr/share/doc/` or `/usr/share/licenses`.

For example, you can find more information about the component `jemalloc-3.6.0` by running the following command:

```
1 ls -l /usr/share/doc/jemalloc-3.6.0/
2
3 total 120
4 -rw-r--r--. 1 root root 1703 Mar 31 2014 COPYING
5 -rw-r--r--. 1 root root 109739 Mar 31 2014 jemalloc.html
6 -rw-r--r--. 1 root root 1084 Mar 31 2014 README
7 -rw-r--r--. 1 root root 50 Mar 31 2014 VERSION
```

However, for some components distributed by CentOS, the license text is not installed in the XenServer product. To view the license text for these components, you can look inside the source RPMs.

Download the source RPMs

We make the source RPMs for the XenServer host available in the following locations:

- For the base ISOs that are periodically released, source files are provided on the [XenServer downloads page](#).
- For updates, the source files are uploaded to a CDN alongside the product RPMs. To download the source files, complete the following steps:

1. Log in to the console of the pool coordinator host.
2. Get the UUIDs of the repositories in use by the pool, by running the following command:

```
1 xe pool-param-get param-name=repositories uuid=<POOL_UUID>
```

3. Check if a proxy is configured and get its URL, by running the following command:

```
1 xe pool-param-get param-name=repository-proxy-url uuid=<POOL_UUID>
```

You cannot use these steps to download the source files through a proxy that has a user name and password configured.

4. If a proxy is configured, in yum configure the proxy for the repositories, by running the following command for each repository:

```
1 yum-config-manager --save --setopt=remote-<REPO_UUID>-source.proxy=<PROXY_URL>
```

5. To download the source RPM of a package, run the following command:

```
1 yumdownloader --disablerepo=* --enablerepo=remote-<REPO_UUID_0>-source,remote-<REPO_UUID_1>-source --source <PKG_NAME>
```

The name of the source file for a specific component is given by the value of “Source RPM” in the detailed information output. For example:

```
1 Source RPM : m4-1.4.16-10.el7.src.rpm
```

Other XenServer components

Supplemental Packs

Supplemental packs are installed into the XenServer host. If you have supplemental packs installed in your host, their RPM information is included when you complete the steps in the previous section of this article.

The source files for supplemental packs are also provided on the [XenServer downloads page](#).

XenCenter

To view information about third-party components included in XenCenter, complete the following steps:

1. In XenCenter, go to **Help > About XenCenter**.
2. Click **View Legal Notices**.

XenServer VM Tools for Windows

The XenServer VM Tools for Windows comprises the following components:

- The Management Agent, which is covered by a proprietary license.
- The Windows I/O drivers, which are covered by the [BSD2 license](#). Copyright Cloud Software Group, Inc.

Licensing information is included in the INF file for each driver. When the drivers are installed on your Windows system by Windows Update or the Management Agent installer, the INF files are stored as `C:\Windows\INF\OEM*.inf`. The Management Agent installer also places the INF files in `C:\Program Files\XenServer\XenTools\Drivers\***.inf`.

Source is not provided for XenServer VM Tools for Windows.

XenServer VM Tools for Linux

The XenServer VM Tools for Linux are covered by the [BSD2 license](#). Copyright Cloud Software Group, Inc.

The archive file provided on the [product download page](#) contains the license file and source files for the tools.

Virtual Appliances

The following virtual appliances are provided as optional components for your XenServer environment:

- Conversion Manager Virtual Appliance
- Workload Balancing Virtual Appliance

These virtual appliances are also CentOS based. You can use the same commands as those given for the XenServer host to get overview and detailed information about the open source packages included in the virtual appliances.

In the console of the virtual appliance, run the following commands:

- For overview information: `rpm -qa --qf '%{ name } -%{ version } : %{ license } \n'`
- For detailed information: `rpm -qai | sed '/^Name /i\\n'`

In addition, the Conversion Manager virtual appliance and Workload Balancing virtual appliance dynamically use some third-party components.

- For Conversion Manager virtual appliance, the license files for these components are located at the following path: `/opt/vpaxcm/conversion`.
- For Workload Balancing virtual appliance, the license files for these components are located at the following path: `/opt/vpx/wlb`.

Source files for the virtual appliances are provided on the [XenServer downloads page](#).

Develop for XenServer®

June 23, 2026

XenServer provides an SDK that enables you to develop or script clients that work with XenServer in the following languages: [PowerShell](#), [Python](#), [C#](#), [C](#), [Go](#), and [Java](#).

To get started developing for XenServer, see our [SDK guide](#).

If you already know what you are doing and want to look up a class or method, see our API reference for [XenServer 8.4](#) or [XenServer 9](#).

Other development tools

To develop an incremental backup solution, see [Changed Block Tracking](#).

Develop drivers and supplemental packs

To develop drivers, supplemental packs, or other software that requires changes to XenServer, contact your Citrix Ready® Partner contact for further information. If you are not yet a Citrix Ready Partner, learn more at [Citrix Ready Partners](#).

Data Governance

June 26, 2026

This article provides information regarding the collection, storage, and retention of logs by XenServer.

XenServer is a server virtualization platform that enables the customer to create and manage a deployment of virtual machines. XenCenter® is the management UI for XenServer. XenServer and XenCenter can collect and store customer data as part of providing the following capabilities:

- **XenServer telemetry** - The telemetry functionality transmits basic licensing information about a XenServer pool. XenServer collects this basic licensing data as necessary for its legitimate interests, including license compliance.
- **Citrix® License telemetry** - XenServer sends telemetry information to your Citrix License Server.
- **Server status reports** - A server status report can be generated on-demand and uploaded to Citrix Insight Services or provided to Support. The server status report contains information that can aid in diagnosing issues in your environment.
- **Automatic updates for the Management Agent** - The Management Agent runs within VMs hosted on a XenServer host or pool. If the server or pool is licensed, the Management Agent can check for and apply updates to itself and to the I/O drivers in the VM. As part of checking for updates, the automatic update feature makes a web request to Cloud Software Group that can identify the VM where the Management Agent runs.
- **XenCenter email alerts** XenCenter can be configured to send email notifications when alert thresholds are exceeded. To send these email alerts, XenCenter collects and stores the target email address.

Telemetry information received by Cloud Software Group is treated in accordance with our [Agreements](#).

XenServer telemetry

The XenServer telemetry functionality collects basic licensing information about your XenServer pools.

When you install XenServer, your pool coordinator gathers telemetry data and uploads it weekly to a Microsoft Azure Cloud environment in the United States. This data does not identify individuals or customers and is sent securely over HTTPS on port 443 to <https://telemetry.ops.xenserver.com/>. No information other than the four elements identified in the following table is collected or transmitted.

Access to this data is restricted to members of the XenServer Operations and Product Management teams.

Telemetry information received by Cloud Software Group is treated in accordance with our [Agreements](#).

Telemetry collected

For each XenServer pool, the pool coordinator collects the following data:

Data collected	Description
UUID	A random unique ID for the telemetry data of this pool. This UUID is not the same as the pool UUID or any other existing identifier. It is not collected in server status reports.
Product version	The version of XenServer installed in this pool.
Sockets (per host)	The number of sockets this host has.
Edition (per host)	The type of license on this host.

This data does not identify individuals or customers and contains no personally identifiable information.

Viewing the telemetry data

The data that XenServer submits is logged on your pool coordinator in `/var/telemetry/telemetry.data`. This file is not collected in the server status logs.

Citrix License telemetry

XenServer also sends telemetry information to your Citrix License Server. For more information, see [Citrix License telemetry](#).

This information is uploaded daily to your Citrix License Server. The data is sent securely over HTTPS on port 8083. This telemetry is required for cloud-based licensing.

Telemetry information received by Cloud Software Group is treated in accordance with our [Agreements](#).

Information collected

For each XenServer pool, the pool coordinator collects the following data:

Data collected	Description
UUID	The pool UUID.
ID	An ID for each instance of telemetry uploaded.
Timestamp	The date and time when the telemetry is uploaded.
Product version	The version of XenServer installed in this pool.
VMs	The number of VMs in the pool.
VDA VM counts	Information about the types of VDA session, the version, and how many VMs are in each category.
Operating Systems	The number of VMs installed with each operating system.
Update source	The source that the pool is configured to receive updates from.
HA enabled	Whether the high availability feature is in use.
Workload Balancing enabled	Whether the Workload Balancing feature is in use.
Conversion Manager installed	Whether the Conversion Manager is installed in the pool.
GPU usage	Information about what GPU and vGPU modes are used and by how many VMs. The VMs are categorized by operating system.
Anti-affinity group count	The number of anti-affinity VM groups used in the pool.
vGPU types	The vGPU types configured in the pool, collected from all vGPU objects and deduplicated.
Corosync version	The version of the Corosync cluster stack, or an empty value if GFS2 clustering is not enabled.
DR enabled	Indicates whether Disaster Recovery (DR) is enabled for the pool.
User Agents connecting to XenServer API	List of user-agent types and versions (XenCenter, SDKs, scripts, tools) that connected to the pool. Helps understand how XenServer is managed.
Windows PV Tools versions in use	Versions of Windows PV Tools used across all Windows VMs in the pool.
Number of VMs using Intellicache	Number of VMs for which Intellicache caching is enabled on their VDIs.

Data collected	Description
Number of MCS-provisioned VMs	Number of VMs provisioned using Citrix MCS.
Number of PVS-provisioned VMs	Number of VMs provisioned using Citrix PVS.
Number of VMs using PVS Accelerator	Number of VMs using PVS Accelerator functionality.
PVS Accelerator enabled status	Indicates whether PVS Accelerator is configured in the pool (true/false).
Number of PXE-booted VMs	Approximate number of VMs configured to boot using PXE/network boot.
Number of SRs with read-caching disabled	Number of Storage Repositories where read-caching has been explicitly disabled.
Storage Repository (SR) types in the pool	List of SR types present anywhere in the pool (for example, lvmoiscsi, gfs2).
Largest VM memory size	Memory (MB) of the largest VM in the pool.
Largest VM vCPU count	vCPU number of the largest VM in the pool.
Number of PCI passthrough VMs (other-config)	Number of VMs configured with PCI passthrough.
Number of SR-IOV VMs	Number of VMs configured to use SR-IOV networking.
Number of USB-passthrough VMs	Number of VMs that have virtual USB devices attached.
vGPU migration usage count	Number of vGPU migration events in the past day.
VM boot mode distribution	Mapping of VM boot modes (BIOS / UEFI / SecureBoot) and number of VMs using each.
Number of VMs with vTPM	Number of VMs in the pool with a virtual TPM device attached.
Number of VMs using DMC	Number of VMs using Dynamic Memory Control. Excludes Dom0, templates, and snapshots.
Number of disks over 2TB	Number of VM disks with virtual_size greater than 2 TB.
Number of VMs with total disks over 2TB	Number of VMs whose total attached disk capacity exceeds 2 TB.
Largest disk size	Size in bytes of the largest disk attached in the pool
VDI count per SR	Mapping of SR type to VDI count, showing how VM storage is distributed across storage repository types.

Data collected	Description
Number of thin provisioned SRs	Number of LVM-based SAN SRs using thin provisioning.
External auth enabled	Indicates whether LDAP external authentication is enabled for the pool.
Secure LDAP enabled	Indicates whether Secure LDAP (LDAPS) is enabled for the pool.

For each host in the pool, the pool coordinator collects the following data:

Data collected	Description
Host UUID (per host)	The host UUID.
System manufacturer	The manufacturer of the server hardware for the host.
System model	The model of the server hardware for the host.
Product version (per host)	The version of XenServer installed on the host.
Edition (per host)	The type of license on this host.
Sockets (per host)	The number of sockets this host has.
Cores (per host)	The number of CPU cores this host has.
Host memory (per host)	The total memory (in bytes) in the host.
dom0 memory (per host)	The amount of memory (in bytes) allocated to the control domain (dom0).
Network mode (per host)	The networking stack used (currently only OpenvSwitch supported).
Last updated time (per host)	When updates were last applied to this host.
Last updated hash (per host)	The identifying checksum for the last set of updates applied to this host.
Boot mode	The host boot mode: BIOS, UEFI, or Secure Boot.
Hardware inventory	Information about the host's server model, NICs, GPUs, and storage controllers, including firmware and driver versions.
Physical CPU cores	The number of physical CPU cores installed on the host.
Hyperthreading status	Indicates whether CPU hyperthreading is enabled on the host.

Data collected	Description
Xen additional parameters	Additional Xen hypervisor command-line parameters beyond the defaults.
Kernel additional parameters	Additional Linux kernel command-line parameters beyond the defaults.
Last host reboot time	Epoch timestamp of the most recent host reboot.
Loaded kernel modules	List of kernel modules currently loaded on the host.
Boot-from-SAN mode	SAN boot mode being used (for example, iBFT for software iSCSI boot).
Number of NUMA nodes	Number of CPU NUMA nodes present on the host.
NRPE enabled status	Indicates if NRPE monitoring is enabled on the host.
SNMP enabled status	Indicates if SNMP monitoring is enabled on the host.
Storage Repository types on host	List of Storage Repository (SR) types attached to this host.
GPU devices present	List of GPU device models detected on the host.
CPU model(s)	CPU model(s) detected on the host (usually one entry).
Multipathing enabled status	Indicates whether multipathing is enabled on this host.
Xen hypervisor version	Version of the Xen hypervisor running on this host.
Applied Xen livepatches	List of Xen livepatches applied to this host.
dom0 Kernel version	Version of the host dom0 kernel.
Applied kernel livepatches	List of kernel livepatches applied to this host.
Detected PCI devices	List of PCI identifiers for hardware components present on the host.
Network bonding modes	List of network bonding modes configured on the host (for example, lacp, balance-slb).
LAS Activation ID (LPEid)	Unique LAS-generated activation ID, reported only for LAS-activated XenServer pools.
Detected storage devices	List of storage vendors and models.

Server status reports

During operation a XenServer host collects and logs various information on the server where XenServer is installed. These logs can be collected as part of a server status report.

A server status report can be generated on-demand. You can upload these reports to Citrix Insight Services or provide them to Support. The server status report contains information that can aid in diagnosing issues in your environment.

Server status reports that are uploaded to Citrix Insight Services are stored in Amazon S3 environments in the United States.

XenServer and XenCenter collect information from the following data sources:

- XenCenter
- XenServer hosts and pools
- Hosted VMs

You can select which data items are included in the server status reports. You can also delete any server status reports that are uploaded to your MyCitrix account on Citrix Insight Services.

Citrix Insight Services does not implement an automatic data retention for server status reports uploaded by the customer. The customer determines the data retention policy. You can choose to delete any server status reports that are uploaded to your MyCitrix account on Citrix Insight Services.

Data collected

A server status report can contain the following log files:

Log type	Contains PII?
<code>device-model</code>	yes
<code>fcoe</code>	yes
<code>firstboot</code>	yes
<code>network-status</code>	yes
<code>process-list</code>	yes
<code>xapi</code>	yes
<code>xenserver-databases</code>	yes
<code>control-slice</code>	maybe
<code>disk-info</code>	maybe

Log type	Contains PII?
hardware-info	maybe
high-availability	maybe
host-crashdump-logs	maybe
kernel-info	maybe
loopback-devices	maybe
message- switch	maybe
multipath	maybe
system-logs	maybe
v6d	maybe
xapi-clusterd	maybe
xapi-debug	maybe
xcp-rrdd-plugins	maybe
xen-info	maybe
xenopsd	maybe
xenserver-config	maybe
xenserver-install	maybe
xenserver-logs	maybe
xha-liveset	maybe
yum	if customized
network-config	if customized
cron	if customized
blobs	no
block-scheduler	no
boot-loader	no
conntest	no
CVSM	no
pam	no
system-services	no
tapdisk-logs	no

Log type	Contains PII?
<code>VM-snapshot-schedule</code>	no
<code>xapi-subprocess</code>	no
<code>xen-bugtool</code>	no
<code>xenserver-domains</code>	no

Management Agent automatic updates

The Management Agent runs within VMs hosted on a XenServer host or pool. If the host or pool is licensed, the Management Agent can check for and apply updates to itself and to the I/O drivers in the VM. As part of checking for updates, the automatic update feature makes a web request to us that can identify the VM where the Management Agent runs.

The web logs captured from the requests made by the Management Agent automatic updates feature are in a Microsoft Azure Cloud environment in the United States. These logs are then copied to a log management server in the United Kingdom.

The web requests made by the Management Agent automatic updates feature are made over HTTPS. Web log files are transmitted securely to the log management server.

You can select whether your VM uses the Management Agent automatic update feature. If you choose to use the Management Agent automatic update feature, you can also choose whether the web request includes the VM identifying information.

Web logs containing information from web requests made by the Management Agent automatic updates feature and the XenCenter check for updates feature can be retained indefinitely.

Data collected

The Management Agent automatic updates web requests can contain the following data points:

Data collected	Description	What we use it for
IP address	The IP address of the VM where the Management Agent is installed	

Data collected	Description	What we use it for
Partial VM UUID	The first four characters of the unique user ID for the VM where the Management Agent is installed	

XenCenter email alerts

XenCenter can be configured to send email notifications when alert thresholds are exceeded. To send these email alerts, XenCenter collects and stores the target email address.

The email address that XenCenter uses to send email alerts is stored on the machine where you installed XenCenter.

You can delete email alerts configured in XenCenter to remove the stored email information.

XenCenter retains the email information used to provide email alerts for the lifetime of the email notification. When you delete the configured email alert, the data is removed.

Data collected

To provide email alerts XenCenter stores the following data points:

Data collected	Description	What we use it for
Email address	The email address for alerts	To send alert and notification emails to
SMTP server	The SMTP server to use	To route the email alerts to the recipient



© 2025 Cloud Software Group, Inc. All rights reserved. This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of Cloud Software Group, Inc. This and other products of Cloud Software Group may be covered by registered patents. For details, please refer to the Virtual Patent Marking document located at <https://www.cloud.com/legal>. XenServer, the XenServer logo, the XenServer 'X' logo, Xen, and the Xen 'X' logo and other marks appearing herein are either registered trademarks or trademarks of Cloud Software Group, Inc. and/or its subsidiaries in the United States and/or other countries. Other marks are the property of their respective owner(s) and are mentioned for identification purposes only. Please refer to Cloud SG's Trademark Guidelines and Third Party Trademark Notices (<https://www.cloud.com/legal>) for more information.