



Citrix Hypervisor 8.2

Contents

About Citrix Hypervisor 8.2 Cumulative Update 1	5
What's new	6
Cumulative Update 1	6
Fixed issues	13
8.2 LTSR (Initial release)	15
Fixed issues	21
Known issues	24
Deprecation	33
System requirements	41
Hardware drivers	46
Configuration limits	49
Guest operating system support	53
Connectivity requirements	57
Quick start	61
Technical overview	84
Technical FAQs	91
Licensing	108
Install	116
Installation and deployment scenarios	124
Upgrade from an existing version	130
Update your hosts	140
Troubleshoot the installation	154
Boot from SAN environments	155

Network boot installations	160
Host partition layout	169
Install on small devices	170
Hosts and resource pools	171
Clustered pools	193
Troubleshoot clustered pools	200
Manage users	206
Role-based access control	214
RBAC roles and permissions	215
Use RBAC with the CLI	226
Networking	230
Manage networking	247
Troubleshoot networking	269
Storage	273
Create a storage repository	279
Thin-provisioned shared GFS2 block storage	300
Manage storage repositories	312
Storage multipathing	323
IntelliCache	330
Storage read caching	334
PVS-Accelerator	338
Graphics overview	347
Prepare host for graphics	351
Create vGPU enabled VMs	361

Memory usage	367
Monitor and manage your deployment	370
Monitor CPU usage	392
Manage virtual machines	396
Windows VMs	400
XenServer VM Tools for Windows	412
Linux VMs	428
Migrate VMs	441
Import and export VMs	446
Delete VMs	461
vApps	462
VM memory	466
Demo Linux virtual appliance	472
Advanced notes for virtual machines	474
Enable VNC for Linux VMs	483
Troubleshoot VM problems	496
High availability	505
Disaster recovery and backup	513
Enable Disaster Recovery	516
vApps	521
Back up and restore hosts and VMs	522
VM snapshots	527
Cope with machine failures	533
Troubleshooting	537

Workload Balancing	540
What's new in Workload Balancing	542
Get started with Workload Balancing	545
Workload Balancing basic tasks	556
Configure Workload Balancing behavior	574
Administer Workload Balancing	603
Certificates for Workload Balancing	632
Troubleshoot Workload Balancing	640
XenServer Conversion Manager	647
What's new in XenServer Conversion Manager	651
Get started with XenServer Conversion Manager	651
Troubleshoot XenServer Conversion Manager	664
Command-line interface	666
Common criteria	791
Third party notices	791
Citrix Hypervisor Open Source Licensing and Attribution	792
Develop for XenServer	797
Data Governance	797

About Citrix Hypervisor 8.2 Cumulative Update 1

April 8, 2024

Citrix Hypervisor is a high-performance hypervisor optimized for virtual app and desktop workloads and based on the Xen Project hypervisor.

Citrix Hypervisor 8.2 Cumulative Update 1 is a long-term-service release (LTSR) of Citrix Hypervisor. It is supported until June 25, 2025.

This documentation reflects features and configurations in the Citrix Hypervisor 8.2 Cumulative Update 1 release:

- To learn about Citrix Hypervisor 8.2 Cumulative Update 1 features, see [What's New](#).
- To download Citrix Hypervisor 8.2 Cumulative Update 1, go to [the Citrix Hypervisor download page](#).
- To get started with Citrix Hypervisor 8.2 Cumulative Update 1, see [Quick Start](#).

The other articles for the latest release of Citrix Hypervisor are listed in the table of contents on the left. If you are on mobile, you can access this table of contents from the menu icon (three horizontal bars) at the top of the page.

XenServer 8

The latest version of Citrix Hypervisor is [XenServer 8](#). If you want to use our [latest release](#) and try out the [new features](#), give XenServer 8 a try. You can use all of its features in a reduced size pool with Trial Edition.

XenServer 8 receives frequent updates that deliver new features, improvements, and bug fixes. If you prefer a more stable feature set, you can stay with Citrix Hypervisor 8.2 Cumulative Update 1 until XenServer 8 becomes an LTSR.

Citrix Hypervisor 8.2 Cumulative Update 1 is supported until June 25, 2025. Plan for the upgrade to XenServer 8 before this date.

Earlier releases

The Citrix Hypervisor product lifecycle strategy for Long Term Service Releases is described in [Lifecycle Milestones for Citrix Hypervisor](#).

About Citrix Hypervisor

Citrix Hypervisor is the complete server virtualization platform from Citrix. The Citrix Hypervisor package contains all you need to create and manage a deployment of virtual x86 computers running on Xen. Citrix Hypervisor is optimized for both Windows and Linux virtual servers.

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

A virtual machine (VM) is 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.

Citrix Hypervisor lets you create VMs, take VM disk snapshots, and manage VM workloads. For a comprehensive list of major Citrix Hypervisor features, visit <https://www.citrix.com/products/citrix-hypervisor>.

XenCenter

XenCenter is a Windows GUI client that provides a rich user experience when managing multiple Citrix Hypervisor servers and resource pools, and the virtual infrastructure associated with them. For more information, see the [XenCenter documentation](#).

What's new

December 7, 2022

About this release

About [Cumulative Update 1 \(CU1\)](#)

About [8.2 LTSR \(initial release\)](#)

Cumulative Update 1

July 15, 2024

About this release

Citrix Hypervisor 8.2 CU1 is the first Cumulative Update for the Citrix Hypervisor 8.2 Long Term Service Release (LTSR). This article provides important information about the Citrix Hypervisor 8.2 CU1 release.

Citrix Hypervisor 8.2 CU1 is available in two commercial editions:

- Premium Edition
- Standard Edition

Citrix Hypervisor 8.2 CU1 and its subsequent hotfixes are available only to customers with Customer Success Services.

To receive these hotfixes through XenCenter, you must also install the latest version of XenCenter and obtain a client ID. For more information, see [Authenticating your XenCenter to receive updates](#).

Support and maintenance

For Citrix Hypervisor 8.2 LTSR, support is available until Jun 25, 2025 on the latest cumulative update.

The Citrix Hypervisor 8.2 initial release is supported until Jun 12, 2022, which is six months following the release of Citrix Hypervisor 8.2 CU1. This support includes troubleshooting of issues and, where feasible, resolution of issues through applying configuration changes or existing product updates or upgrades.

For Citrix Hypervisor 8.2 LTSR, maintenance is available until Jun 25, 2025 on the latest cumulative update. This maintenance includes both critical and functional hotfixes.

Until Jun 12, 2022, we issue critical hotfixes for both Citrix Hypervisor 8.2 initial release and Citrix Hypervisor 8.2 CU1. We issue any functional hotfixes for Citrix Hypervisor 8.2 CU1 only.

For more information about support and maintenance, see [Product Lifecycle Support Policy](#).

Included in Citrix Hypervisor 8.2 CU1

Citrix Hypervisor 8.2 CU1 rolls-up all previously issued Citrix Hypervisor 8.2 hotfixes and simultaneously introduces new fixes for issues reported on Citrix Hypervisor 8.2. For more information, see [Fixed issues in Citrix Hypervisor 8.2 CU1](#).

Some performance or non-functional improvements are also included. For more information, see [Improvements in Citrix Hypervisor 8.2 CU1](#).

To minimize changes within the LTSR product, no additional features are included in Citrix Hypervisor 8.2 CU1.

Improvements in Citrix Hypervisor 8.2 CU1

In addition to the rolled-up hotfixes, Citrix Hypervisor 8.2 CU1 includes some performance or non-functional improvements that are available for all licensed LTSR customers.

XenServer Conversion Manager 8.3.1

XenServer Conversion Manager 8.3.1 - the latest version of the XenServer Conversion Manager virtual appliance - allows you to convert VMs in parallel, enabling you to migrate your entire VMware environment to XenServer quickly and efficiently. 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.

For more information on how to use the XenServer Conversion Manager, see [XenServer Conversion Manager](#).

Citrix VM Tools for Linux 8.2.1-1

Citrix VM Tools for Linux 8.2.1-1 adds support for Red Hat Enterprise Linux 9.

Note:

You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 9 VMs as Red Hat does not support memory ballooning with the Xen hypervisor.

For more information, see [Install Citrix VM Tools for Linux](#).

Support for kdump and kexec in Linux VMs

The [kdump](#) utility and [kexec](#) command are now supported on Linux VMs running on Citrix Hypervisor.

Cipher suite changes

To improve security, weaker cipher suites have been removed from the list of cipher suites that are supported for SSH communication. For information about the cipher suites that are now supported, see [Communicate with Citrix Hypervisor servers and resource pools](#).

Improved jumbo frames support

Jumbo frames are now supported for traffic on VM networks. For more information, see [Jumbo frames](#).

Authenticated download of updates

To provide a more secure service for hotfix downloads, XenCenter now requires that you authenticate it with Citrix to automatically download and apply hotfixes.

To receive these hotfixes through XenCenter, you must also install the latest version of XenCenter and obtain a client ID JSON file. For more information, see [Authenticating your XenCenter to receive updates](#).

Automated Updates in XenCenter

Automated Updates were previously restricted to Citrix Hypervisor Premium Edition customers or Citrix Virtual Apps and Desktops customers. However, in pools with hotfix XS82ECU1053 applied, this feature is available to all users. For more information, see [Authenticating your XenCenter to receive updates](#).

XenCenter PuTTY update

The version of PuTTY embedded in XenCenter 8.2.4 and later is updated to 0.76.

Changes to guest operating system support

For the full list of supported guest operating systems in Citrix Hypervisor 8.2 Cumulative Update 1, see [Guest operating system support](#).

Important:

Support for all paravirtualized (PV) VMs was removed in Citrix Hypervisor 8.1. However, in Citrix Hypervisor 8.2 Cumulative Update 1 and later, unsupported 32-bit PV VMs are prevented from starting. Ensure that you do not have any 32-bit PV VMs in your environment before upgrading to Citrix Hypervisor 8.2 CU 1.

Added Citrix Hypervisor 8.2 CU1 now supports the following new guests:

- Windows Server 2022 (with the latest XenServer VM Tools for Windows)

Note:

At release of Citrix Hypervisor 8.2 Cumulative Update 1, Windows Server 2022 has not been validated on this platform with the Server Virtualization Validation Program (SVVP).

Windows Server 2022 is supported for production use only with XenServer VM Tools for Windows version 9.2.1 or later installed. You can get the latest version of the XenServer VM Tools for Windows from the [Citrix Hypervisor Product Download](#) page

- Rocky Linux 8
- Gooroom 2

From hotfix XS82ECU1036 and onwards

- Debian Bullseye 11
- Ubuntu 22.04

From hotfix XS82ECU1046 and onwards

- Red Hat Enterprise Linux 9 (with the latest Citrix VM Tools for Linux)

From hotfix XS82ECU1050 and onwards

- Debian Bookworm 12 (with the latest Citrix VM Tools for Linux)
- Rocky Linux 9 (with the latest Citrix VM Tools for Linux)
- CentOS Stream 9 (with the latest Citrix VM Tools for Linux)

From hotfix XS82ECU1069 and onwards

- Ubuntu 24.04

Removed Citrix Hypervisor 8.2 CU1 no longer supports the following guests:

- Windows Server 2012
- Windows 8.1
- Debian Jessie 8
- Debian Stretch 9
- Ubuntu 16.04
- SUSE Linux Enterprise Desktop 12 SP3, 12 SP4
- SUSE Linux Enterprise Desktop 15 SP3
- SUSE Linux Enterprise Server 12 SP3

- CentOS 8
- Windows Server 2012 R2
- Ubuntu 18.04
- SUSE Linux Enterprise Server 12 SP4

Support for new processors

The following processors are now supported in Citrix Hypervisor 8.2 CU1:

- Intel® Xeon® 84xx/64xx/54xx/44xx/34xx (Sapphire Rapids)
 - If you are installing on Sapphire Rapids hardware with four or more sockets, ensure that you use a base installation ISO that was released after Jun 1, 2023. The latest version of the ISO is available at the [Citrix Hypervisor Product Download](#) page.
 - If you are installing on an HPE ProLiant Gen11 server with Sapphire Rapids hardware, ensure that you use a base installation ISO that was released after Jun 1, 2023. The latest version of the ISO is available at the [Citrix Hypervisor Product Download](#) page.
 - If you are installing on Sapphire Rapids hardware not covered by the previous bullets, you can use any version of the ISO. However, if you use a version of the ISO that was released before Jun 1, 2023, ensure that you apply hotfix [XS82ECU1026](#).
- Intel® Xeon® E-23xx and E-23xxG Processors (Rocket Lake)
- AMD EPYC™ 9xx4 (Zen4/Genoa)

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

Compatibility with Citrix Hypervisor 8.2 virtual appliances

You can use the Workload Balancing 8.3.0 and XenServer Conversion Manager 8.3.1 virtual appliances with your Citrix Hypervisor 8.2 CU1 pool. To use the latest virtual appliances with Citrix Hypervisor 8.2 CU1, ensure that you use the latest version of XenCenter provided with Citrix Hypervisor 8.2.

For information about the latest version of Workload Balancing, see [What's new in Workload Balancing](#). For information about the XenServer Conversion Manager, see [XenServer Conversion Manager](#).

Installation options

Citrix Hypervisor 8.2 CU1 is available to download from the [Citrix Hypervisor Product Download](#) page in the following package:

- **Citrix Hypervisor 8.2 Cumulative Update 1 Base Installation ISO** comprises both a base Citrix Hypervisor 8.2 installation and the fixes that make up the cumulative update. Use this ISO to create a fresh installation of Citrix Hypervisor 8.2 including CU1.

Notes:

Before beginning a fresh installation, review the [System Requirements](#) and [Installation](#).

After installation of Citrix Hypervisor 8.2 CU1, the internal product version number is shown as 8.2.1.

Previously, an update ISO was provided for update from Citrix Hypervisor 8.2 to Citrix Hypervisor 8.2 Cumulative Update 1. Citrix Hypervisor 8.2 is no longer supported and the update ISO has been moved to the *Prior Versions* section of the downloads page.

Upgrade and update options

You cannot upgrade or update to Citrix Hypervisor 8.2 CU1 from earlier versions that are now out of support. Instead you must create a fresh installation. For more information, see [Installing](#).

Optional components updated in Citrix Hypervisor 8.2 CU1

The following are new/updated optional components for the Citrix Hypervisor 8.2 CU1 release. All other optional components remain the same as those available with the Citrix Hypervisor 8.2 release.

- [XenCenter 2024.2.0](#)
- Software Development Kit 8.2.3

Licensing

Upgrade your Citrix License Server to version 11.16 or higher in order to use all Citrix Hypervisor 8.2 licensed features.

Note:

When upgrading to Citrix Hypervisor 8.2 Cumulative Update 1, we recommend that you transition from using Citrix License Server virtual appliance to using the Windows-based Citrix License Server. The Citrix License Server virtual appliance is no longer supported.

For more information about Citrix Hypervisor 8.2 licensing, see [Licensing Overview](#).

Interoperability with Citrix products

Citrix Hypervisor 8.2 CU1 is interoperable with the following Citrix Virtual Apps and Desktops versions:

- Citrix Virtual Apps and Desktops 2203
- Citrix Virtual Apps and Desktops 2112
- Citrix Virtual Apps and Desktops 1912 LTSR
- Citrix Virtual Apps and Desktops orchestrated by using Citrix Cloud.

We recommend that you use this Citrix Hypervisor LTSR with a Citrix Virtual Apps and Desktops LTSR.

Citrix Hypervisor 8.2 CU1 is also supported with Citrix Provisioning 2112 and 1912 LTSR.

Localization support

The localized versions of XenCenter (Simplified Chinese and Japanese) are also available in this release.

Product documentation

To access Citrix Hypervisor 8.2 LTSR product documentation, see [Citrix Hypervisor Product Documentation](#).

Fixed issues

January 9, 2023

Rolled up Hotfixes in CU1

Citrix Hypervisor 8.2 Cumulative Update 1 includes the following Citrix Hypervisor 8.2 hotfixes:

- XS82E001 - <https://support.citrix.com/article/CTX277444>
- XS82E002 - <https://support.citrix.com/article/CTX285938>
- XS82E003 - <https://support.citrix.com/article/CTX280214>
- XS82E004 - <https://support.citrix.com/article/CTX284749>
- XS82E005 - Superseded
- XS82E006 - <https://support.citrix.com/article/CTX285536>
- XS82E007 - <https://support.citrix.com/article/CTX283446>

- XS82E008 - <https://support.citrix.com/article/CTX283510>
- XS82E009 - <https://support.citrix.com/article/CTX283516>
- XS82E010 - <https://support.citrix.com/article/CTX285172>
- XS82E011 - <https://support.citrix.com/article/CTX286459>
- XS82E012 - <https://support.citrix.com/article/CTX286796>
- XS82E013 - <https://support.citrix.com/article/CTX286800>
- XS82E014 - <https://support.citrix.com/article/CTX286804>
- XS82E015 - <https://support.citrix.com/article/CTX292897>
- XS82E016 - <https://support.citrix.com/article/CTX292625>
- XS82E017 - <https://support.citrix.com/article/CTX294145>
- XS82E018 - Limited Availability
- XS82E019 - Limited Availability
- XS82E020 - <https://support.citrix.com/article/CTX313808>
- XS82E021 - <https://support.citrix.com/article/CTX306540>
- XS82E022 - <https://support.citrix.com/article/CTX306423>
- XS82E023 - <https://support.citrix.com/article/CTX312232>
- XS82E024 - <https://support.citrix.com/article/CTX306481>
- XS82E025 - <https://support.citrix.com/article/CTX310674>
- XS82E026 - <https://support.citrix.com/article/CTX313807>
- XS82E028 - <https://support.citrix.com/article/CTX318325>
- XS82E029 - <https://support.citrix.com/article/CTX319717>
- XS82E030 - <https://support.citrix.com/article/CTX322578>
- XS82E031 - <https://support.citrix.com/article/CTX327907>
- XS82E032 - <https://support.citrix.com/article/CTX324257>
- XS82E033 - <https://support.citrix.com/article/CTX328166>
- XS82E034 - <https://support.citrix.com/article/CTX330706>

Additional issues fixed in CU1

General

- When multiple VMs start at the same time, Workload Balancing recommends balancing the VMs placement on all servers in the pool evenly. However, sometimes Workload Balancing might recommend putting many VMs on the same Citrix Hypervisor server. This issue occurs when Workload Balancing gets late feedback from XAPI about VM placement. (CA-337867)
- A runaway process that logs excessively can fill the log partition and prevent log rotation. This issue is resolved by rotating the log files before they exceed 100MB. (CA-356624)
- When installing Citrix Hypervisor on a system with the HBA355i Adapter Card the system hangs on the install screen. (CA-357134)

- The DNS settings in xsconsole are not retained after host reboot. (CA-355872)
- If an SR scan reports errors during SR attach, the attach process can fail. (CA-355401)
- You cannot use snapshots for VMs located on an NFS SR provided by a Tintri VMstore file system. (CA-359453)
- Under certain conditions the tapdisk process for a VM can crash when updating performance statistics. (CA-355145)

Guests

- On some servers, VMs with GPU/PCI pass-through configured can fail to boot and log the error; “Operation not permitted.”(CA-356386)
- [Fixed in the latest Citrix VM Tools for Linux] If you attempt to install the Citrix VM Tools for Linux on a fully up-to-date CentOS 8 system, you see the error: **Fatal Error: Failed to determine Linux distribution and version**. This is caused by changes in that CentOS 8 updates release on Dec 08, 2020. To work around this issue, specify the OS when installing the Citrix VM Tools for Linux: `./install.sh -d centos -m 8`. However, if you use this workaround, the operating system information is not reported back to the Citrix Hypervisor server and does not appear in XenCenter. (CA-349929)
- The `kdump` utility and `kexec` command are now supported on Linux VMs running on Citrix Hypervisor. (CP-24801)

XenCenter

- If you have FIPS compliance enabled on the system where XenCenter is installed, you cannot import or export VMs in OVF/OVA format or import Virtual Hard Disk images. (CA-340581)
- When using XenCenter to upgrade pools in parallel and apply all released hotfixes after the upgrade, the hotfix files can be downloaded multiple times. This can cause a delay in hotfix application and increase the amount of downloaded data. (CA-359700)

8.2 LTSR (Initial release)

July 15, 2024

About this release

Citrix Hypervisor is a high-performance hypervisor optimized for virtual app and desktop workloads and based on the Xen Project hypervisor.

Citrix Hypervisor 8.2 is a Long Term Service Release which seeks to maximize stability in terms of the feature set.

Citrix Hypervisor 8.2 is available in the following editions:

- Premium Edition
- Standard Edition

For information about the features available in each edition, see the [Citrix Hypervisor Feature Matrix](#).

New features and improvements in Citrix Hypervisor 8.2

Citrix Hypervisor 8.2 introduces enhanced features and functionality for application, desktop, and server virtualization use cases. All Citrix Hypervisor 8.2 features are available to all licensed Citrix Virtual Apps and Desktops customers or Citrix DaaS customers.

Increased configuration limits

The following configuration limits have been increased:

- The maximum host RAM is now 6 TB
- The maximum number of logical processors per host is now 448 CPUs

For more information, see [Configuration limits](#).

Enable and disable read caching from within XenCenter

The read caching feature improves performance on NFS, EXT3/EXT4, SMB, or GFS2 SRs that host multiple VMs cloned from the same source. This feature can now be enabled and disabled for each individual SR from the XenCenter console. You might want to disable read caching in the following cases:

- You have no file-based SRs
- You do not have any cloned VMs
- You have insufficient memory available to allocate to dom0 to derive any performance benefits

For more information, see [Changing SR Properties](#).

Changes to guest operating system support

The set of guest operating systems that Citrix Hypervisor supports has been updated. For more information, see [Guest operating system support](#).

Added Citrix Hypervisor now supports the following additional guest operating systems:

- SUSE Linux Enterprise Server 12 SP5 (64-bit)
- Ubuntu 20.04 (64-bit)
- Gooroom 2 (64-bit) - requires [Hotfix XS82E021 - For Citrix Hypervisor 8.2](#)

Removed

- Windows 7
- Windows Server 2008 SP2
- Windows Server 2008 R2 SP1

Changes to processor support

The following processors are now supported:

- Xeon 83xxH(L)/63xxH(L)/53xxH (Cooper Lake SP)
- Xeon 83xx/63xx/53xx/43xx (Ice Lake SP)
- AMD EPYC 7xx3 Zen3 (Milan)

To get the full benefits of these processors, ensure that you install the latest hotfixes for Citrix Hypervisor 8.2.

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

Security improvements

Install a TLS certificate on your Citrix Hypervisor server Citrix Hypervisor now enables easy installation of server TLS certificates.

The Citrix Hypervisor server comes installed with a default TLS certificate. However, to use HTTPS to secure communication between Citrix Hypervisor and Citrix Virtual Apps and Desktops, you must install a new certificate. The certificate authority issuing the certificate must be trusted by the Citrix Virtual Apps and Desktop installation.

This feature provides a mechanism to update host certificates that does not require the user to have access to the Citrix Hypervisor server file system. It also verifies that the certificate and key files are valid and in the correct format.

You can install a TLS certificate on the Citrix Hypervisor server by using one of the following methods:

- XenCenter. For more information, see [Install a TLS certificate on your server](#) in the XenCenter documentation.
- xe CLI. For more information, see [Install a TLS certificate on your server](#).
- API. For more information, see the [Management API guide](#).

This feature also provides XenCenter alerts when a server TLS certificate is about to expire. For more information, see [System Alerts](#) in the XenCenter documentation.

Enforcing use of the TLS 1.2 protocol Citrix Hypervisor now enforces the use of the TLS 1.2 protocol for any HTTPS traffic between Citrix Hypervisor and an external network. All Citrix Hypervisor components use the TLS 1.2 protocol when communicating with each other.

As part of this feature the legacy SSL mode and support for the TLS 1.0/1.1 protocol have been removed. Ensure you disable legacy SSL mode in your pools before upgrading or updating them to Citrix Hypervisor 8.2. If you have any custom scripts or clients that rely on a different protocol, update these components to use TLS 1.2.

Disaggregation of the XenServer VM Tools

The XenServer VM Tools (formerly Citrix VM Tools) are now provided as two separate components on the [Citrix Hypervisor download page](#):

- XenServer VM Tools for Windows
- Citrix VM Tools for Linux

As a result, the `guest-tools.iso` file has been removed from the Citrix Hypervisor installation.

Providing the tools as a separate component removes the need for hotfixes to apply updates to a tools ISO stored on the Citrix Hypervisor server.

For more information, see [Linux VMs](#) and [Windows VMs](#).

Updates to the Workload Balancing appliance and Conversion Manager appliance

Both of these additional components now include the following improvements:

- Platform updated to CentOS 7.7
- Performance gains from using JSON-RPC to communicate with Citrix Hypervisor
- OpenSSL updated to version 1.1.1
- Other third-party libraries updated for security and performance improvements
- TLS 1.2 is now enforced to ensure security

Installation options

Citrix Hypervisor 8.2 is available to download from the [Citrix Hypervisor Product Download](#) page in the following packages:

- Citrix Hypervisor 8.2 Update ISO. Use this file to apply Citrix Hypervisor 8.2 as an update to Citrix Hypervisor 8.1 or 8.0.
- Citrix Hypervisor 8.2 Base Installation ISO. Use this file to create a fresh installation of Citrix Hypervisor 8.2 or to upgrade from XenServer 7.1 CU2 or 7.0.

Important:

- If you use XenCenter to upgrade your hosts, update your XenCenter installation to the latest version supplied on the Citrix Hypervisor 8.2 download page before beginning.
- Always upgrade the pool master before upgrading any other hosts in a pool.
- Ensure that you update your XenServer 7.1 to Cumulative Update 2 before attempting to upgrade to Citrix Hypervisor 8.2.
- Legacy SSL mode is no longer supported. Disable this mode on all hosts in your pool before attempting to upgrade to the latest version of Citrix Hypervisor. To disable legacy SSL mode, run the following command on your pool master before you begin the upgrade: `xe pool-disable-ssl-legacy uuid=<pool_uuid>`
- The Container Management supplemental pack is no longer supported. After you update or upgrade to the latest version of Citrix Hypervisor, you can no longer use the features of this supplemental pack.
- The vSwitch Controller is no longer supported. Disconnect the vSwitch Controller from your pool before attempting to update or upgrade to the latest version of Citrix Hypervisor.
 1. In the vSwitch controller user interface, go to the **Visibility & Control** tab.
 2. Locate the pool to disconnect in the **All Resource Pools** table. The pools in the table are listed using the IP address of the pool master.
 3. Click the cog icon and select **Remove Pool**.
 4. Click **Remove** to confirm.

After the update or upgrade, the following configuration changes take place:

- Cross-server private networks revert to single-server private networks.
- Any Quality of Service settings made through the DVSC console are no longer applied. Network rate limits are no longer enforced.
- ACL rules are removed. All traffic from VMs is allowed.
- Port mirroring (RSPAN) is disabled.

After update or upgrade, if you find any leftover state about the vSwitch Controller in your pool, clear the state with the following CLI command: `xe pool-set-vswitch-controller address=`

Before you start

Before beginning installation or migrating from an older version, review the following articles:

- [System requirements](#)
- [Known issues](#)
- [Deprecations and removals](#)

For information about the installation, upgrade, or update process, see [Install](#)

Licensing

Customers must upgrade their Citrix License Server to version 11.16 or higher to use all Citrix Hypervisor 8.2 licensed features.

For more information about Citrix Hypervisor 8.2 licensing, see [Licensing](#).

Hardware compatibility

For the most recent additions and advice for all hardware compatibility questions, see the Citrix Hypervisor [Hardware Compatibility List](#).

If you have VMs with attached virtual GPUs, ensure that supported drivers are available before upgrading to the latest release of Citrix Hypervisor. For more information, see both the [Hardware Compatibility List](#) and the GPU vendor documentation.

Interoperability with Citrix products

Citrix Hypervisor 8.2 is interoperable with Citrix Virtual Apps and Desktops 7.15 LTSR, 1912 LTSR, and 2006.

Citrix Hypervisor 8.2 is interoperable with Citrix Provisioning 7.15 LTSR, 1912 LTSR, and 2006.

Citrix Hypervisor 8.2 is interoperable with Citrix Cloud.

Localization support

The localized version of XenCenter (Simplified Chinese and Japanese) is also available in this release. In previous releases, the localized version of XenCenter was provided as a separate component. In Citrix Hypervisor 8.2 and later, all localized version of XenCenter are contained in the same `.msi` installation file as the English version.

Product documentation

To access Citrix Hypervisor 8.2 product documentation, see [Citrix Hypervisor 8.2 Product Documentation](#).

To access the latest XenCenter product documentation, see [XenCenter Product Documentation](#).

Documentation can be updated or changed after the initial release. We suggest that you subscribe to the [Document History](#) RSS feed to learn about updates.

Fixed issues

January 9, 2023

This article lists issues present in previous releases that are fixed in this release.

General

- On Citrix Hypervisor 8.1 systems that were updated from Citrix Hypervisor 8.0, after XAPI restart you might see errors from the following services:
 - `usb-scan.service`
 - `storage-init.service`
 - `xapi-domains.service`
 - `mpathcount.service`
 - `create-guest-templates.service`

The errors in these services can present various different issues, for example, VMs that have been configured to start automatically might not start. The cause of this issue is `xapi-wait-init-complete.service` not being enabled. (CA-333953)

- Improvements to boot time, memory accounting, and stability of Citrix Hypervisor on systems with large amount of RAM. (CP-33195)

- A system with a software FCoE connection might experience a persistent memory leak in the Dom0 kernel. This memory leak eventually results in a host crash, sudden loss of connectivity, or other issue. (CA-332618)
- Some of the methods that take a `DateTime` parameter in the C# SDK and the corresponding PowerShell module cmdlets fail with an internal error. (CA-333871)
- USB pass-through does not work with version 2.00 devices with a speed less than or equal to 12 Mbps. (CA-328130)
- In rare error conditions, the XAPI process on the pool master can leak file descriptors for stunnel connections. This issue can cause the pool to become non-operational. (CA-337546)
- If there is malformed or unexpected content in the `/etc/passwd` or `/etc/group` file on your server, an upgrade from an earlier version of XenServer to Citrix Hypervisor 8.1 can fail. (CA-336696)
- When shutting down the system, a crash might occur and the system is rebooted instead. (CA-334114)
- If you have previously attempted to install the incorrect version of a driver disk, subsequent driver disk installations can fail because of data remaining in the yum cache. (CA-330961)
- Sometimes the RRD metrics related to VBDs that Workload Balancing collects from the Citrix Hypervisor control domain can be in an incorrect format. Workload Balancing now ignores incorrectly formatted metrics, collects the metrics again, and sends a WARNING log. (CA-335950)

Guests

- When installing Citrix VM Tools on a Windows VM, the installation might fail with the following error message: `Service 'Citrix XenServer Windows Management Agent' (XenSvc) could not be installed. Verify that you have sufficient privileges to install system services.` Work around this issue by uninstalling the Citrix VM Tools, rebooting the VM, and then installing the new tools.
- On CentOS 8 VMs with the Citrix VM Tools installed, boot times can be slow. (CA-333687)
- A VM migration from a pool member that takes more than 12 hours can fail with a connection reset error. This failure is caused by an idle connection between the pool master and the pool member timing out. (CA-333610)
- If the Management Agent is installed on your Windows VM, attempting to copy more than 1 MB of text to the clipboard can cause your VM to become unresponsive. (CA-326354)
- When objects such as SRs are destroyed, their RRDs are not removed from memory, which can cause memory usage to grow over time. (CA-325582)

Storage

- When running Reclaim Space on a thinly provisioned LUN with more than 2 TB of free space, the operation fails with an `ioctl not supported` error. (CA-332782)
- Creating a VDI with Unicode characters in either the name or description causes the database backup script to fail with an error on a GFS2 SR. (CA-335367)
- The `xcp-rrdd-iostat` daemon does not recognize VDIs associated with IntelliCache as valid, causing spam in its log file: `Could not file device with physical path...` (CA-144246)

XenCenter

- When creating an LVM SR from XenCenter and passing CHAP credentials, the operation might fail with an authentication error. (CA-337280)

Improvements

- Greater tolerance for I/O during the leaf coalesces. This change benefits to customers who are taking regular snapshots on active VMs. (CP-32204)
- Greater tolerance for coalescing large leaves. This change benefits customers who have fast storage. (CP-32204)
- The xe CLI client that can be installed on a remote Windows or Linux system contains the following improvements:
 - The xe CLI client can now only upload configuration files that are less than 32 MiB
 - The xe CLI client only uploads or downloads files listed in the original command line arguments
 - Diagnostics xe CLI commands are limited to users with the Pool Admin or Pool Operator role

Update the version of the xe CLI on your remote systems to the latest version provided with Citrix Hypervisor 8.2.

- The read and write latency per device metrics give us the average latency per operation. Previously, this average was taken over all operations ever performed. The average is now taken over the preceding five seconds. This change fixes an issue where the metrics showed constant read or write latencies when no disk operations were in progress. (CA-336067)

Known issues

July 15, 2024

This article contains advisories and minor issues in the Citrix Hypervisor 8.2 release and any workarounds that you can apply.

General

- The backup and restore capability in xsconsole is temporarily not available. We are actively investigating a replacement solution. (CP-48776)
- If you host the Citrix License Server virtual appliance version 11.14 or earlier on your Citrix Hypervisor server, you see a warning when upgrading or updating to Citrix Hypervisor 8.2 Cumulative Update 1. The warning states that this virtual appliance is a PV VM that is no longer supported. We recommend that you transition to using the Windows-based Citrix License Server before upgrading or updating.

The latest version of Citrix License Server is available from the [Citrix Licensing download page](#).

- 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 power off and then start the VM. Rebooting the VM, for example, by clicking 'Reboot' in XenCenter, does not update the VM's feature set. (CA-188042)
- The increase in the amount memory allocated to dom0 in Citrix Hypervisor 8.0 can mean there is slightly less memory available for running VMs. On some hardware, you cannot run the same number of VMs with Citrix Hypervisor 8.2 as you can on the same hardware with XenServer 7.6 and earlier. (CP-29627)
- When attempting to use the serial console to connect to a Citrix Hypervisor server, 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)
- When read caching is enabled, it is slower to read from the parent snapshot than from the leaf. (CP-32853)
- After upgrading to Citrix Hypervisor 8.2, when there is a lot of VM activity on a server in a pool that uses NFS storage, connections through ENIC to external storage can become temporarily blocked (5–35 minutes). VMs on that server can freeze and their consoles can become unresponsive. Attempts to ping in or out of the affected server subnet fail during these times. To fix this

issue, install version 4.0.0.8-802.24 or later of the enic driver. For more information, see [Driver Disk for Cisco enic 4.0.0.11 - For Citrix Hypervisor 8.x CR](#). (XSI-916)

- When attempting to log in to the dom0 console with an incorrect password, you receive the following error message: `When trying to update a password, this return status indicates that the value provided as the current password is not correct`. This error message is expected even though it relates to a password change, not a login. Try to log in with the correct password.
- If an Active Directory user inherits the pool admin role from an AD group that has spaces in its name, the user cannot log in to Citrix Hypervisor 8.2 CU1 through SSH.

To work around this issue, you can do one of the following:

- Rename this group in AD to remove the spaces from its name.
- Create a new group in AD with the pool admin role and without spaces in its name. Add the user to this new group.

In XenCenter, ensure that your new or renamed group is listed in the **Users** tab, either by adding or by removing and re-adding the group. Ensure that the user is associated with the new or renamed group. (CA-363207)

- In clustered pools, a network outage might cause the following issues: inability to reconnect to the GFS2 storage after a host reboot, inability to add or remove hosts in a pool, difficulties in managing the pool. If you experience any of these problems in your clustered pool, contact Citrix Support for advice about recovering your environment. (XSI-1386)

Graphics

- When you start in parallel many VMs with AMD MxGPU devices attached, some VMs might fail with a VIDEO_TDR_FAILURE. This behavior might be due to a hardware limitation. (CA-305555)
- When NVIDIA T4 added in pass-through mode to a VM on some specific server hardware, that VM might not power on. (CA-360450)
- On hardware with NVIDIA A16/A2 graphics cards, VMs with vGPUs can sometimes fail to migrate with the internal error “Gpumon_interface.Gpumon_error([S(Internal_error);S((Failure “No vGPU available”))])”. To recover the VM from this state, you must shut it down and start it again. (CA-374118)
- [Fixed in [XS82ECU1031](#)] When running a large number of virtual GPU or GPU pass-through enabled VMs, rebooting some VMs on the host might cause a performance glitch in the other VMs.

Guests

- Citrix Hypervisor no longer supports VMs that run in PV mode. In previous releases, after upgrading your Citrix Hypervisor server to the latest version, these unsupported VMs might still run. However, in Citrix Hypervisor 8.2 CU1, 32-bit PV mode VMs no longer start-up. Ensure that you remove any PV VMs from your pool or convert these VMs to HVM mode before upgrading to your pool to Citrix Hypervisor 8.2 CU1. For more information, see [Upgrade from PV to HVM guests](#). (CP-38086)
- If you attempt to revert a VM to a scheduled VSS snapshot that was created with Citrix Hypervisor 8.0 or earlier, the VM does not boot. The boot fails with the following error: `This operation cannot be performed because the specified virtual disk could not be found`. This failure is because the VSS snapshot capability has been removed from Citrix Hypervisor 8.1 and later. (CA-329469)

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)
- The guest UEFI boot capability provided in Citrix Hypervisor 8.0 was an experimental feature. Citrix Hypervisor 8.2 does not support migrating UEFI boot VMs created in Citrix Hypervisor 8.0 to Citrix Hypervisor 8.2. Shut down UEFI boot VMs before upgrading to Citrix Hypervisor 8.2 from Citrix Hypervisor 8.0. (CA-330871)
- When UEFI-boot VMs start, they show a TianoCore logo. (CP-30146)
- On Windows VMs, when updating the `xenbus` driver to version 9.1.0.4, ensure that you complete both of the requested VM restarts. If both restarts are not completed, the VM might revert to emulated network adapters and use different settings, such as DHCP or different static IP addressing.

To complete the second restart, you might be required to use a local account to log into the Windows VM. When you log in, you are prompted to restart.

If you are unable to log in to the Windows VM after the first restart, you can use XenCenter to restart the VM and complete the `xenbus` driver installation. (CP-34181)

- 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:

- In the UEFI console, type the following commands.

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

- 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)

- On a Windows VM, after you have installed the version 9.x XenServer VM Tools for Windows, you might see both the previous and the latest version of the tools or management agent listed in your **Installed Programs**:
 - (PREVIOUS) Citrix XenServer Windows Management Agent
 - (LATEST) Citrix Hypervisor PV Tools

The previous version of the management agent is not active and does not interfere with the operation of the latest version. We advise that you do not manually uninstall **Citrix XenServer Windows Management Agent** because this can disable the [xenbus](#) driver and cause the VM to revert to emulated devices.

- On all Windows versions, Windows Update treats the latest xennet driver as a Manual Update and does not install it automatically. You can check the status of your driver installations by going to **Windows Settings > Update & Security > View Update History > Driver Updates**. If this occurs, you can install the driver by downloading the XenServer VM Tools for Windows from the [Citrix Hypervisor download page](#), and manually installing the driver inside the MSI file. (CA-350838)
- On a Windows VM, 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 VMs Service Manager. (CA-340227)
- On a Windows VM with more than 8 vCPUs, Receive Side Scaling might not work because the xenvif driver fails to set up the indirection table. (CA-355277)
- When attempting to update a Windows 10 VM from 1909 to 20H2 or later, the update might fail with a blue screen showing the error: INACCESSIBLE BOOT DEVICE.

To make it less likely that this failure occurs, you can take the following steps before attempting to update:

1. Update the XenServer VM Tools for Windows on your VM to the latest version.
2. Snapshot the VM.
3. In the VM registry, delete the following values from the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Hardware Profiles\Current\HwProfile\Parameters\Boot\Device key: ActiveDeviceID, ActiveInstanceID, and ActiveLocationInformation

This issue is still under investigation. (XSI-1075)

- When creating a Windows VM from a template that is set to not automatically update its drivers, the created VM is incorrectly set to update its drivers. To work around this issue, run the following command: `xe pool-param-set policy-no-vendor-device=true uuid=<pool-uuid>`. This command ensures that future VMs created from the template are correctly set to not automatically update drivers. VMs that were previously generated from the template are not changed. (CA-371529)

Linux guests

- You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Rocky Linux 8, Rocky Linux 9, or CentOS Stream 9 VMs as these operating systems do not support memory ballooning with the Xen hypervisor. (CA-378797) (CP-45141)
- 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)
- The Citrix VM Tools for Linux can provide an incorrect value for the free memory of the VM that is higher than the correct value. (CA-352996)
- If you install Debian 10 (Buster) by using PXE network boot, do not add `console=tty0` to the boot parameters. This parameter can cause issues with the installation process. Use only `console=hvc0` in the boot parameters. (CA-329015)
- After a CentOS 8 VM with only one CPU is migrated to a new Citrix Hypervisor server, the first time a CPU-bound command runs on the VM, it times out. To work around this issue, you can assign more than one CPU to the VM and restart it. (XSI-864)
- If you attempt to shut down a CentOS Stream 9 VM by using XenCenter or the xe CLI, the shut-down process pauses and times out after 1200s. This behaviour is caused by a kernel issue in kernel-5.14.0-362.el9.

The following workarounds are available:

- To work around a single instance of the issue, you can shut down the VM from inside the guest operating system.
- To prevent the issue from occurring for your VM, downgrade your VM to use kernel-5.14.0-354.el9 by running the following commands in the VM:

```
1 # yum install kernel-5.14.0-354.el9
2 # grubby --set-default /boot/vmlinuz-5.14.0-354.el9.x86_64
```

- When installing Debian 11 32-bit on a VM using a QEMU emulated network device, the installation might fail. This issue is caused by the Xen PV drivers being missing from the installer kernel.

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.

We recommend to work around this issue by disabling Wayland or upgrading your Mesa drivers:

- 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.

- Update your Mesa drivers

Run the following commands:

```
1 sudo add-apt-repository ppa:kisak/kisak-mesa
2 sudo apt update
3 sudo apt upgrade
```

Restart the VM to complete the process.

Installation

- If you are using a legacy disk layout, the control domain has less space available to it than the current layout (4 GB vs 18 GB).

In this case, when attempting to apply the Citrix Hypervisor 8.2 or Citrix Hypervisor 8.2 Cumulative Update 1 update to an earlier, you might receive the error message “the server does not have enough space”. This error happens because installation of the Citrix Hypervisor update requires sufficient free space to avoid filling the disk, which is not possible with the legacy layout.

If you receive this error, you cannot update to Citrix Hypervisor 8.2 or Citrix Hypervisor 8.2 Cumulative Update 1. Do a fresh installation instead. (CA-268846)

- If any vSwitch Controller state remains in your pool after an update or upgrade, clear the state with the following CLI commands:

```
1 xe pool-set-vswitch-controller address=
2 xe pool-param-set uuid=<uuid> vswitch-controller=
```

(CA-339411)

- Occasionally, booting a Citrix Hypervisor server from FCoE SAN using software FCoE stack can cause the host to stop responding. This issue is caused by a temporary link disruption in the host initialization phase. If the host fails to respond for a long time, you can restart the host to work around this issue.
- When upgrading to or installing Citrix Hypervisor 8.2 from an ISO located on an IIS server, the install or upgrade can fail and leave your servers unable to restart. The remote console shows the GRUB error: “File ‘/boot/grub/i3860pc/normal.mod’ not found. Entering rescue mode”. This issue is caused by the IIS configuration causing package files to be missing. To work around this issue, ensure that double escaping is allowed on IIS before extracting the installation ISO on it. (XSI-1063)
- [Fixed in XenCenter 8.2.6] If you have the Container Supplemental Pack installed on your XenServer 7.1 CU2 host and attempt to upgrade to Citrix Hypervisor 8.2 CU1 by using XenCenter, you are prevented from upgrading because the supplemental pack is no longer supported. For more information, see [Deprecation](#).

To work around the issue, use the xe CLI to complete the upgrade. This upgrade removes the Container Management Supplemental Pack. For more information, see [Upgrade Citrix Hypervisor servers by using the xe CLI](#). (XSI-1250)

- [Fixed in the reissued Base Installation ISO - Feb 24, 2022] If your server has a NIC that requires the ice driver, this NIC is not available as part of the Citrix Hypervisor 8.2 CU1 installation process, as such you cannot configure it as the management interface or use it to retrieve installation files from a network server. Instead, use a different NIC during the installation process. (CA-363735)
- [Fixed in the reissued Base Installation ISO - Jun 5, 2022] If you attempt to install Intel® Xeon® 84xx/64xx/54xx/44xx/34xx (Sapphire Rapids) hardware with four or more sockets by using an installation ISO released before Jun 1, 2023, the installation fails. To successfully install on this hardware, use a version of the base installation ISO released after Jun 1, 2023.

Internationalization

- Non-ASCII characters, for example, characters with accents, cannot be used in the host console. (CA-40845)
- In a Windows VM with XenServer VM Tools for Windows installed, copy and paste of double-byte characters can fail if using the default desktop console in XenCenter. The pasted characters appear as question marks (?).

To work around this issue, you can use the remote desktop console instead. (CA-281807)

Storage

- If you use GFS2 SRs and have two servers in your clustered pool, your cluster can lose quorum and fence during an upgrade. To avoid this situation, either add a server to or remove a server from your cluster. Ensure that you have either one or three servers in your pool during the upgrade process. (CA-313222)
- If you are using a GFS2 SR, ensure that you enable storage multipathing for maximum resiliency. If storage multipathing is not enabled, file system block writes might not fully complete in a timely manner. (CA-312678)
- Citrix Hypervisor doesn't support MCS full clone VMs with GFS2 SRs. (XSI-832)
- If a Citrix Virtual Apps and Desktops or Citrix DaaS deployment with its VMs hosted on Citrix Hypervisor 8.2 uses multiple GFS2 SRs in a single MCS catalog, VMs in the catalog cannot access the VDIs during deployment. The error "VDI is currently in use" is reported. (XSI-802)
- If you are using a GFS2 SR and your cluster network is on a non-management VLAN, you cannot add or remove hosts to your clustered pool. (XSI-1604)
- If you use HPE 3PAR hardware for your storage repository and, with a previous release of XenServer, you use ALUA1 for your host persona, when you upgrade to Citrix Hypervisor 8.2 multipathing no longer works. To work around this issue, migrate your host persona to ALUA2. For more information, see https://support.hpe.com/hpsc/doc/public/display?docId=emr_na-c02663749&docLocale=en_US.
- After removing an HBA LUN from a SAN, you might see log messages and I/O failures when querying Logical Volume information. To work around this issue, reboot the Citrix Hypervisor server. (XSI-984)
- You cannot set or change the name of the tmpfs SR used by the PVS Accelerator Supplemental Pack. When the `type` is `tmpfs`, the command `xe sr-create` disregards the value set for `name-label` and instead uses a fixed value. If you attempt to run the command `xe sr-param-set` to change the name of the tmpfs SR, you receive the error `SCRIPT_MISSING`.
- You cannot run more than 200 PVS-Accelerator-enabled VMs on a Citrix Hypervisor server. (CA-365079)
- If you install hotfix [XS82ECU1065](#), you cannot attach or reattach a read-only NFS SR.
- Sometimes, failures when leaf coalescing can lead to the allocated space being reported as higher than the expected value. This issue is caused by temporary snapshots being left behind by the failed leaf coalesce. To work around this issue, either shut down the VM or reduce its I/O load. The storage garbage collector can then succeed in coalescing the leaf. (XSI-1517)

Third-party

- 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.

Workload Balancing

- During the Workload Balancing maintenance window, Workload Balancing is unable to provide placement recommendations. When this situation occurs, you see the error: “4010 Pool discovery has not been completed. Using original algorithm.”The Workload Balancing maintenance window is less than 20 minutes long and by default is scheduled at midnight. (CA-359926)
- In XenCenter, the date range showed on the Workload Balancing Pool Audit Report is incorrect. (CA-357115)
- For a Workload Balancing virtual appliance version 8.2.2 and later that doesn’t use LVM, you cannot extend the available disk space. (CA-358817)

XenCenter

- Changing the font size or dpi on the computer on which XenCenter is running can result in the user interface appearing incorrectly. The default font size is 96 dpi; Windows 8 and Windows 10 refer to this font size as 100%. (CA-45514) (CAR-1940)
- On Windows 10 (1903 and later) VMs, there can be a delay of a few minutes after installing the XenServer VM Tools before the **Switch to Remote Desktop** option is available in XenCenter. You can restart the toolstack to make this option appear immediately. (CA-322672)
- It is not advisable to update the same pool from concurrent instances of XenCenter because this might disrupt the update process.

If more than one instance of XenCenter is attempting to install multiple hotfixes on a pool, a server might fail to install a hotfix with the error: “The update has already been applied to this server. The server will be skipped.”This error causes the whole update process to stop. (CA-359814)

To workaround this issue:

1. Ensure that no other XenCenter instance is in the process of updating the pool
2. Refresh the update list in the **Notifications > Updates** panel
3. Start the update from the beginning

- In XenCenter 8.2.3, importing an OVF or OVA file can be significantly slower than in earlier versions of XenCenter. (CP-38523)
- In XenCenter, when you attempt to import an OVF package or a disk image from a folder containing a hash character (#) in its name, the import fails with a null reference exception. (CA-368918)
- XenCenter shows a message marking Dynamic Memory Control (DMC) as deprecated. This deprecation is no longer the case. DMC is supported in future releases.
- In pools with hotfix XS82ECU1029 applied that have GFS2 SRs, using XenCenter to generate a server status report (SSR) can fail. To work around this issue, generate your SSRs by running the following command in the host console: `xenserver-status-report`. (CA-375900)
- If a hotfix requires another hotfix to already be installed as a prerequisite, XenCenter does not display the name of the prerequisite hotfix. You can find the prerequisite information in the article on <https://support.citrix.com> for the hotfix you are trying to install. (CA-383054)

Deprecation

July 15, 2024

The announcements in this article give you advanced notice of platforms, Citrix products, and features that are being phased out so that you can make timely business decisions. Citrix monitors 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 [Product Lifecycle Support Policy](#) article.

Deprecations and removals

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

Deprecated items are not removed immediately. Citrix continues 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 Citrix Hypervisor.

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

Item	Deprecation announced in	Removed in	Alternative
VM Recovery Mode	XenCenter 2024.2.1	XenCenter 2024.2.1	
PuTTY installed with XenCenter	XenCenter 8.2.6	XenCenter 2024.2.0 (for Citrix Hypervisor 8.2 CU1)	With future versions of XenCenter, you must install your own instance of PuTTY or OpenSSH on the system where XenCenter is installed.
Standard Edition license	June 3, 2024	Standard Edition licenses issued before this date are supported until their expiry.	For ways to get a license, see the XenServer website .
Support for Bromium Secure Platform in Windows VMs	8.2 CU1	8.2 CU1	
Health Check	XenCenter 8.2.6	XenCenter 8.2.7	
Upload of server status reports to Citrix Insight Services (CIS) from XenCenter	XenCenter 8.2.6	XenCenter 8.2.7	Upload your status reports manually through the CIS website.
PuTTY installed with XenCenter	XenCenter 8.2.6		With future versions of XenCenter, you must install your own instance of PuTTY or OpenSSH on the system where XenCenter is installed.
Localization of XenCenter into Japanese and Simplified Chinese	XenCenter 8.2.6		
Software Fibre Channel over Ethernet (FCoE)	8.2 CU1		

Item	Deprecation announced in	Removed in	Alternative
Support for Debian Bullseye 11 32-bit Intel GVT-g	8.2 CU1		
Separate PVS-Accelerator Supplemental Pack	8.2 CU1		In future versions of XenServer, these capabilities will be included in the core product.
Support for the following Linux operating systems: CoreOS, Ubuntu 16.04, Ubuntu 18.04, Debian Jessie 8, Debian Stretch 9, SUSE Linux Enterprise Server 12 SP3, SUSE Linux Enterprise Server 12 SP4, SUSE Linux Enterprise Desktop 12 SP3, SUSE Linux Enterprise Desktop 12 SP4, CentOS 8	8.2 CU1	8.2 CU1	Upgrade your VMs to a later version of their operating system where available.
Support for Windows Server 2012 R2	8.2 CU1		Upgrade your VMs to a later version of their operating system.
Support for Windows Server 2012 and Windows 8.1	8.2 CU1	8.2 CU1	Upgrade your VMs to a later version of their operating system.

Item	Deprecation announced in	Removed in	Alternative
Transfer VM	8.2 CU1	8.2 CU1 (XenCenter 8.2.3)	Use the latest release of XenCenter. Since XenCenter 8.2.3, the mechanism used for OVF/OVA import/export and single disk image import has been simplified and these operations are now performed without using the Transfer VM.
Measured Boot Supplemental Pack	8.2 CU1	8.2 CU1	
Container Management Supplemental Pack	8.2	8.2	
Support for Hewlett-Packard Integrated Lights-Out (iLO)	8.2	8.2	
Support for the following legacy processors: Xeon E3/5/7 family - Sandy Bridge , Xeon E3/5/7 v2 family - Ivy Bridge	8.2	8.2	
The guest-tools.iso file included in the Citrix Hypervisor installation ISO	8.2	8.2	Download the Citrix VM Tools for Windows or for Linux from the Citrix Hypervisor downloads page .

Item	Deprecation announced in	Removed in	Alternative
Support for Windows 7, Windows Server 2008 SP2, and Windows Server 2008 R2 SP1	8.2	8.2	Upgrade your VMs to a later version of their operating system.
Legacy SSL mode and support for the TLS 1.0/1.1 protocol	8.2	8.2	
Cross-server private networks	8.2	8.2	
The following xe CLI log commands: <code>diagnostic-db-log</code> , <code>log-set-output</code> , <code>log-get-keys</code> , <code>log-get</code> , <code>log-reopen</code>	8.2		
The vSwitch Controller (see Notes)	8.1	8.2	
Legacy partition layouts: DOS partition layout, Old GPT partition layout. Note, this change also removes support for servers with less than 46 GB of primary disk space.	8.1		
VSS and quiesced snapshots	8.1	8.1	
Support for Ubuntu 14.04	8.1	8.1	

Item	Deprecation announced in	Removed in	Alternative
Support for all paravirtualized (PV) VMs.	8.1	8.1	Upgrade your VMs to a later version of their operating system before moving to the latest version of Citrix Hypervisor. For more information, see Upgrade from PV to HVM guests .
Legacy drivers: qla4xxx , qla3xxx , netxen_nic , qlge , qlcnict	8.0		
XenCenter installer bundled with the Citrix Hypervisor installation media.	8.0	8.0	Download the XenCenter installer from the Downloads page instead.
XenCenter connections to XenServer hosts that are version 6.x and earlier.	8.0	8.0	Upgrade your out-of-support XenServer hosts.
Support for Nutanix integration.	8.0	8.0	

Item	Deprecation announced in	Removed in	Alternative
Support for the following legacy processors: Opteron 13xx Budapest , Opteron 23xx/83xx Barcelona , Opteron 23xx/83xx Shanghai , Opteron 24xx/84xx Istanbul , Opteron 41xx Lisbon , Opteron 61xx Magny Cours , Xeon 53xx Clovertown , Xeon 54xx Harpertown , Xeon 55xx Nehalem , Xeon 56xx Westmere-EP , Xeon 65xx/75xx Nehalem-EX , Xeon 73xx Tigerton , Xeon 74xx Dunnington	8.0	8.0	For information about supported processors, see the Hardware Compatibility List .

Item	Deprecation announced in	Removed in	Alternative
Support for <code>qemu-trad</code> . It is no longer possible to use <code>qemu-trad</code> by setting <code>platform-device-model=qemu-trad</code> . All VMs created with <code>qemu-trad</code> device profile get automatically upgraded to <code>qemu-upstream-compat</code> profile.	8.0	8.0	
Support for the following guest templates: Debian 6 Squeeze, Ubuntu 12.04, Legacy Windows, Asianux Server 4.2, 4.4, and 4.5, NeoKylin Linux Security OS 5, Linx Linux 6, Linx Linux 8, GreatTurbo Enterprise Server 12, Yinhe Kylin 4	8.0	8.0	
Legacy Windows drivers from the Citrix VM Tools ISO	8.0	8.0	

Notes

Health Check

Logs for the Health Check service are retained by Windows for troubleshooting purposes. To remove these logs, delete them manually from `%SystemRoot%\System32\Winevt\Logs` on the Win-

dows machine running XenCenter.

Dynamic Memory Control (DMC)

This feature was previously listed as deprecated. The deprecation notice was removed on Jan 30, 2023. DMC is supported in future releases of XenServer.

Disconnecting the vSwitch controller

The vSwitch Controller is no longer supported. Disconnect the vSwitch Controller from your pool before attempting to update or upgrade to the latest version of Citrix Hypervisor.

1. In the vSwitch controller user interface, go to the **Visibility & Control** tab.
2. Locate the pool to disconnect in the **All Resource Pools** table. The pools in the table are listed using the IP address of the pool master.
3. Click the cog icon and select **Remove Pool**.
4. Click **Remove** to confirm.

After the update or upgrade, the following configuration changes take place:

- Cross-server private networks revert to single-server private networks.
- Any Quality of Service settings made through the DVSC console are no longer applied. Network rate limits are no longer enforced.
- ACL rules are removed. All traffic from VMs is allowed.
- Port mirroring (RSPAN) is disabled.

After update or upgrade, if you find any leftover state about the vSwitch Controller in your pool, clear the state with the following CLI command: `xe pool-set-vswitch-controller address=`

System requirements

July 15, 2024

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

Warning:

Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor CLI is also installed. A standalone remote Citrix Hypervisor CLI can be installed on any RPM-based Linux distribution. For more information, see [Command-line interface](#).

Citrix Hypervisor server system requirements

Although Citrix Hypervisor is usually deployed on server-class hardware, Citrix Hypervisor 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 Citrix Hypervisor hardware specifications.

The Citrix Hypervisor server must be a 64-bit x86 server-class machine devoted to hosting VMs. Citrix Hypervisor 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.

Citrix Hypervisor can use:

- Up to 6 TB of RAM
- Up to 16 physical NICs
- Up to 448 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 Citrix Hypervisor server are:

CPUs

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

To support VMs running Windows or more recent versions of Linux, you require an Intel VT or AMD-V 64-bit x86-based system with one or more CPUs.

Note:

To run Windows VMs or more recent versions of Linux, enable hardware support for virtualization on the Citrix Hypervisor server. Virtualization support is an option in the BIOS. It is possible that your BIOS might have virtualization support disabled. For more information, see your BIOS documentation.

To support VMs running supported paravirtualized Linux, you require a standard 64-bit x86-based system with one or more CPUs.

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.

Citrix Hypervisor requires an IPv4 network for management and storage traffic.

Notes:

- Ensure that the time setting in the BIOS of 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 Citrix Hypervisor configuration, we recommend that you configure serial console access. For hosts that do not have physical serial port or where suitable physical infrastruc-

ture is not available, investigate whether you can configure an embedded management device. For example, Dell DRAC. For more information about setting up serial console access, see [CTX228930 - How to configure console access on XenServer or Citrix Hypervisor](#).

XenCenter system requirements

XenCenter has the following system requirements:

- **Operating System:**

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

- **.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

XenCenter is compatible with all supported versions of Citrix Hypervisor.

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 servers, up to a maximum of 64. Before you create a pool or join a server to an existing pool, ensure that all servers in the pool meet the following requirements.

Hardware requirements

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

- The CPU vendor (Intel, AMD) must be the same on all CPUs on all servers.
- To run HVM virtual machines, all CPUs must have virtualization enabled.

Other requirements

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

- It must have a consistent IP address (a static IP address on the server 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 master (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 server 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 server to the pool. After the server has joined the pool, you can reconfigure the management interface again.
- It must be running the same version of Citrix Hypervisor, at the same patch level, as servers already in the pool.
- It must be configured with the same supplemental packs as the servers already in the pool. Supplemental packs are used to install add-on software into the Citrix Hypervisor control domain, dom0. To prevent an inconsistent user experience across a pool, all servers in the pool must have the same supplemental packs at the same revision installed.
- It must have the same Citrix Hypervisor license as the servers already in the pool. You can change the license of any pool members after joining the pool. The server with the lowest license determines the features available to all members in the pool.

Citrix Hypervisor servers 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](#).

Hardware drivers

July 17, 2024

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 Citrix Hypervisor 8.2 Cumulative Update 1 includes third-party drivers that have been certified as compatible with Citrix Hypervisor. A list of the drivers included in-box with your initial Citrix Hypervisor installation is given in the summary article [Driver versions for XenServer and Citrix Hypervisor](#).

Updates to drivers

We regularly deliver updated versions of these drivers as driver disk ISO files. For example, these updates can be provided to enable new hardware or resolve issues with existing hardware.

Driver disk ISO files are released on the website <https://support.citrix.com>. You can find these updates in the following places:

- The latest update to a driver is listed in the summary article [Driver versions for XenServer and Citrix Hypervisor](#).

Driver Name	Driver Type	Citrix Hypervisor 8.2 Cumulative Update 1 LTSR
aacraid	Raid/Storage Controller	1.2.1[50877]-custom Update available: 1.2.1.60001
arcmsr	Raid/Storage Controller	v1.40.00.09-20180709
be2iscsi	iSCSI/Storage Controller	11.4.0.1
be2net	Ethernet/NIC	12.0.0.0
bfa	Fibre Channel HBA/Storage Controller	3.2.25.1
bnx	Ethernet/NIC	3.2.25.1
bnx2/netxtreme2	Ethernet/NIC	2.2.6a
bnx2fc/netxtreme2	FCoE Driver/	2.12.13 Update available: (Dell hardware) 2.12.20
bnx2i/netxtreme2	iSCSI Driver/	2.11.25.0
bnx2x/netxtreme2	Ethernet/NIC	1.715.8 Update available: (Dell hardware) 1.715.20
bnx2x/netxtreme2	Ethernet/NIC	1.10.0-216.0.119.1 Update available: (Dell hardware) 1.10.0-216.0.119.1

- Use the search on the support.citrix.com site to find updated versions for the driver you need.

Citrix Hypervisor

enic

Q

Filter by

All types

Articles

Security Bulletins

Software Updates

Search results found for “enic”

Sort by Created date

Driver Disk for Cisco enic 4.4.0.1-For Citrix Hypervisor 8.2 LTSR

A driver disk for customers running the Citrix Hypervisor 8.2 LTSR release who use Cisco's enic driver and wish to use the latest version.

Created: 13 Apr 2023 | Citrix Hypervisor, XenServer

Driver Disk for Cisco enic 4.2.0.26-For Citrix Hypervisor 8.2 LTSR

This driver disk is for customers running the Citrix Hypervisor 8.2 LTSR release who use Cisco's enic driver and wish to use the latest version.

Created: 21 Jul 2022 | Citrix Hypervisor

Driver Disk for Cisco enic 4.0.0.11-For Citrix Hypervisor 8.x CR

Driver Disk for Cisco enic 4.0.0.11-For Citrix Hypervisor 8.x CR

Created: 10 Aug 2020 | Citrix Hypervisor

- Subscribe to the Citrix Hypervisor software updates [RSS feed](#) to be informed of new drivers and hotfixes as they are released.

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

Support for drivers

Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor. These drivers are not supported.

Before a driver can be supported with Citrix Hypervisor, 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 Citrix Hypervisor environment and that it is compatible with Citrix Hypervisor 8.2 CU1.

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 Citrix Hypervisor.

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 Citrix Hypervisor 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 Citrix Hypervisor and we publish the driver through <https://support.citrix.com>.

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

Installing driver disks

You can install a driver disk by one of the following methods:

- [During Citrix Hypervisor installation](#)
- [By using XenCenter](#)
- [By using the command line](#)

Configuration limits

July 15, 2024

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

- Virtual machine limits
- Citrix Hypervisor server 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 (see note 1)
Virtual CPUs per VM (Windows)	32
Memory	
RAM per VM	1.5 TiB (see note 2)
Storage	
Virtual Disk Images (VDI) (including CD-ROM) per VM	255 (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 (GFS2)	16 TiB

Item	Limit
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	6

Notes:

1. Consult your guest OS documentation to ensure that you do not exceed the supported limits.
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. Some 32-bit Windows operating systems can support more than 4 GiB of RAM through use of the physical address extension (PAE) mode. For more information, see your guest operating system documentation and [Guest operating system support](#).
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.

Citrix Hypervisor server limits

Item	Limit
Compute	

Item	Limit
Logical processors per host	512 (see note 1)
Concurrent VMs per host	1000 (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

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. For larger pools (over 32 hosts), we recommend allocating at least 8GB RAM to the

Control Domain (Dom0). For systems running over 500 VMs or when using the PVS Accelerator, we recommend allocating at least 16 GB RAM to the Control Domain. For information about configuring Dom0 memory, see [CTX220763 - How to change dom0 Memory](#).

- 3. For NVIDIA vGPU, 128 vGPU accelerated VMs per host with 4xM60 cards (4x32=128 VMs), or 2xM10 cards (2x64=128 VMs). For Intel GVT-g, 7 VMs per host with a 1,024 MB aperture size. Smaller aperture sizes can further restrict the number of GVT-g VMs supported per host. This figure might change. For the current supported limits, see the [Hardware Compatibility List](#).
- 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	64 (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)

Item	Limit
Multipathed LUNs per host (used by storage repositories)	150 (see note 2)
VDIs per SR (NFS, SMB, EXT, GFS2)	20000
VDIs per SR (LVM)	1000
Attached VDIs per SR (all types)	600
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.
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](#).

Guest operating system support

July 15, 2024

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 (32-bit) [Latest tested version is 22H2]	1 GB	4 GB	24 GB (40 GB or more recommended)
Windows 10 (64-bit) [Latest tested version is 22H2]	2 GB	1.5 TB	32 GB (40 GB or more recommended)
Windows Server 2016, Windows Server Core 2016 (64-bit)	1 GB	1.5 TB	32 GB (40 GB or more recommended)
Windows Server 2019, Windows Server Core 2019 (64-bit)	1 GB	1.5 TB	32 GB (40 GB or more recommended)
Windows Server 2022, Windows Server Core 2022 (64-bit)	1 GB	1.5 TB	32 GB (40 GB or more recommended)
CentOS 7 (64-bit)	2 GB	1.5 TB	10 GB
CentOS Stream 9 (64-bit) (see note 3)	2 GB	1.5 TB	10 GB
Red Hat Enterprise Linux 7 (64-bit)	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 2)	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 SP1, 15 SP2, 15 SP3, 15 SP4 (64-bit)	1 GB	1.5 TB	8 GB
SUSE Linux Enterprise Desktop 15 SP4 (64-bit)	1 GB	1.5 TB	8 GB
Oracle Linux 7 (64-bit)	2 GB	1.5 TB	10 GB
Oracle Linux 8 (64-bit)	2 GB	1.5 TB	10 GB

Operating System	Minimum RAM	Maximum RAM	Minimum Disk Space
Scientific Linux 7 (64-bit)	2 GB	1.5 TB	10 GB
Debian Buster 10 (64-bit)	1 GB	1.5 TB	10 GB
Debian Bullseye 11 (32-bit) (see note 1)	512 MB	1.5 TB	10 GB
Debian Bullseye 11 (64-bit) (see note 1)	512 MB	1.5 TB	10 GB
Debian Bookworm 12 (64-bit) (see note 3)	1 GB	1.5 TB	10 GB
Ubuntu 20.04 (64-bit)	512 MB	1.5 TB	10 GB
Ubuntu 22.04 (64-bit) (see note 1)	1 GB	1.5 TB	10 GB
Ubuntu 24.04 (64-bit) (see note 4)	1 GB	1.5 TB	10 GB
NeoKylin Linux Advanced Server 7.2 (64-bit)	1 GB	1.5 TB	10 GB
Gooroom 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 3)	2 GB	1.5 TB	15 GB

Note:

1. Requires hotfix [XS82ECU1036](#).
2. Requires hotfix [XS82ECU1046](#). Customers who wish to use this guest OS must also install Citrix VM Tools for Linux v8.2.1-1 or later, available to download from the [Citrix Hypervisor product download page](#).
3. Requires hotfix [XS82ECU1050](#). Customers who wish to use this guest OS must also install Citrix VM Tools for Linux v8.2.1-1 or later, available to download from the [Citrix Hypervisor product download page](#).
4. Requires hotfix [XS82ECU1069](#). Customers who wish to use this guest OS must also install Citrix VM Tools for Linux v8.2.1-1 or later, available to download from the [Citrix Hypervisor product download page](#).

- Citrix VM Tools for Linux is only supported on the Linux guest operating systems listed above.

- All supported operating systems run in HVM mode.
Support for paravirtualized (PV) VMs was removed in Citrix Hypervisor 8.1. Remove PV VMs from your environment before moving to Citrix Hypervisor 8.2 CU 1.
- 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.

- Some 32-bit Windows operating systems can support more than 4 GB of RAM by using physical address extension (PAE) mode. To reconfigure a VM with greater than 4 GB of RAM, use the `xe` CLI, not XenCenter, as the CLI doesn't impose upper bounds for `memory-static-max`.
- Citrix Hypervisor supports all SKUs (editions) for the listed versions of Windows.
- Windows 11 is not supported on Citrix Hypervisor 8.2 CU1.

Long-term guest support

Citrix Hypervisor 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 Citrix Hypervisor was released, but are now no longer supported by their vendors.

Citrix no longer offers support for these operating systems (even if they remain listed in the table of supported guests or their templates remain available on your Citrix Hypervisor hosts). While attempting to address and resolve a reported issue, Citrix assesses if the issue directly relates to an out-of-support operating system on a VM. To assist in making that determination, Citrix 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, Citrix will not investigate the issue further.

Note:

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

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

Connectivity requirements

July 18, 2024

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.

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. For more information, see Update your XenServer hosts

Domain	Port	Direction	Details
citrix.com and subdomains	443	Outbound	XenCenter accesses subdomains on the citrix.com domain to download hotfixes. For more information, see Update your Citrix Hypervisor hosts
storage.googleapis.com	443	Outbound	XenCenter accesses this domain to download hotfixes. For more information, see Update your Citrix Hypervisor hosts

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
pvupdates.vmd.citrix.com	443	Outbound	The XenServer VM Tools for Windows poll information on this site to see whether updates are available for the management agent.
downloadns.citrix.com.edgesuite.net	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 Citrix Hypervisor

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

Not all ports need to be open, depending on your deployment and requirements.

Source	Destination	Type	Port	Details
Citrix Hypervisor	Citrix Hypervisor	TCP	80, 443	Intra-host communication between members of a resource pool using the management API
	Citrix License Server	TCP	27000	Handles initial connection for license requests
		TCP	7279	Check-in/check-out of licenses
	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

Source	Destination	Type	Port	Details
XenCenter	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
	Clustering	TCP	8892, 8896, 21064	Communication between all pool members in a clustered pool
		UDP	5404, 5405	
	Citrix Hypervisor	TCP	22	SSH
		TCP	443	Management using the management API
		TCP	5900	VNC for Linux VMs
Other clients	Citrix Hypervisor	TCP	3389	RDP for Windows VMs
		TCP	80, 443	Any client that uses the management API to communicate with Citrix Hypervisor servers
		TCP		

Note:

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

Additional port information

- [Active Directory integration](#)
- [Citrix Provisioning Services](#)

Quick start

April 16, 2024

This article steps through how to install and configure Citrix Hypervisor 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 servers, which provides the foundation to migrate running VMs between servers 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 Citrix Hypervisor and XenCenter. It is intended for those users who want to administer Citrix Hypervisor by using XenCenter. For information on how to administer Citrix Hypervisor using the Linux-based `xe` commands through the Citrix Hypervisor Command Line Interface (CLI), see [Command-line interface](#).

Terminology and abbreviations

- *Server*: a physical computer that runs Citrix Hypervisor
- *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 Citrix Hypervisor servers and their VMs
- *Storage Repository (SR)*: a storage container in which virtual disks are stored

Major components

Citrix Hypervisor

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

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

Citrix Hypervisor 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 Citrix Hypervisor servers, 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 Citrix Hypervisor and XenCenter

In this section, you set up a minimum Citrix Hypervisor installation.

What you'll learn

You'll learn how to:

- Install Citrix Hypervisor on a single physical server
- Install XenCenter on a Windows computer
- Connecting XenCenter and Citrix Hypervisor 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 Citrix Hypervisor server

- A Windows computer to run the XenCenter application
- Installation files for Citrix Hypervisor and XenCenter

The Citrix Hypervisor server computer is dedicated entirely to the task of running Citrix Hypervisor 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 [Citrix Hypervisor Downloads](#).

Install the Citrix Hypervisor server

All servers have at least one IP address associated with them. To configure a static IP address for the server (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 Citrix Hypervisor server:

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

Note:

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

2. Back up data you want to preserve. Installing Citrix Hypervisor 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 Citrix Hypervisor** screen, select your keyboard layout for the installation.
7. When the **Welcome to Citrix Hypervisor Setup** screen is displayed, select **Ok**.
8. Read and accept the Citrix Hypervisor 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 Citrix Hypervisor server.
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 server 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 server 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 Citrix Hypervisor**.
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 server.
After the server reboots, Citrix Hypervisor 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 server.

Install XenCenter

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

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.

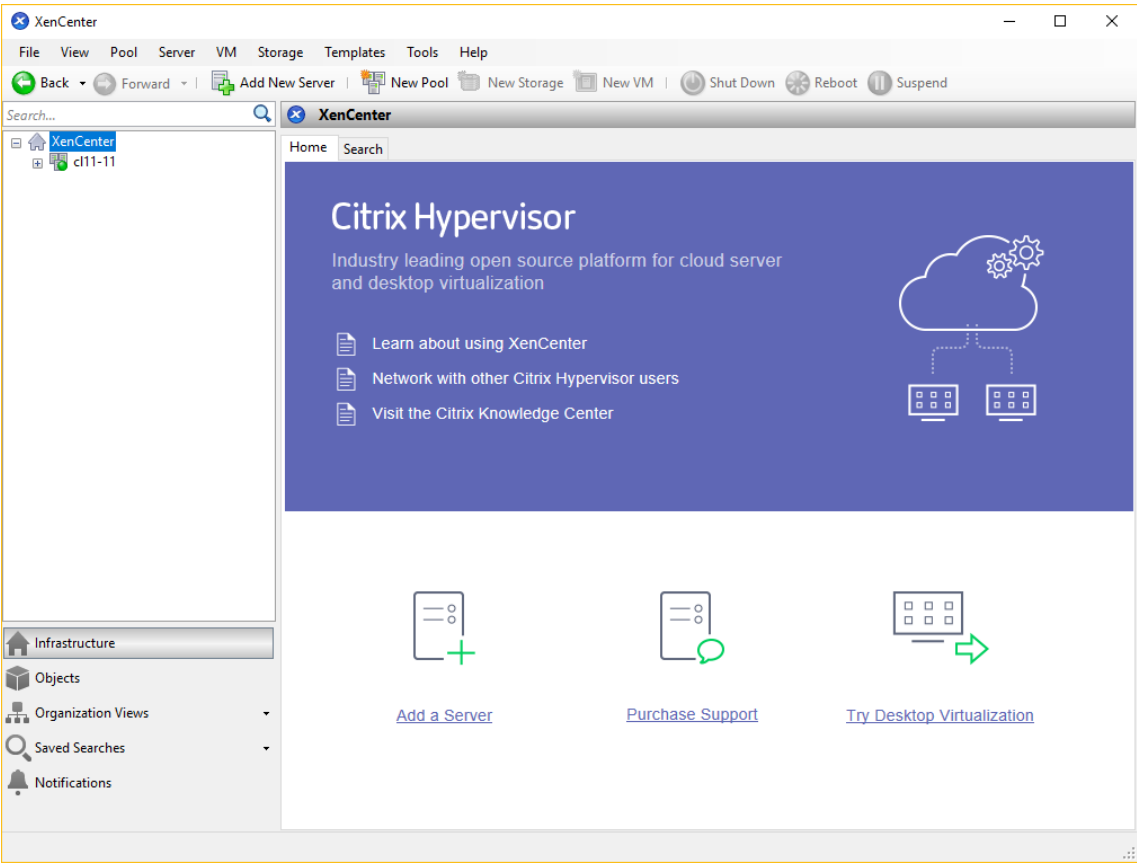
Connect XenCenter to the Citrix Hypervisor server

This procedure enables you to add a server to XenCenter.

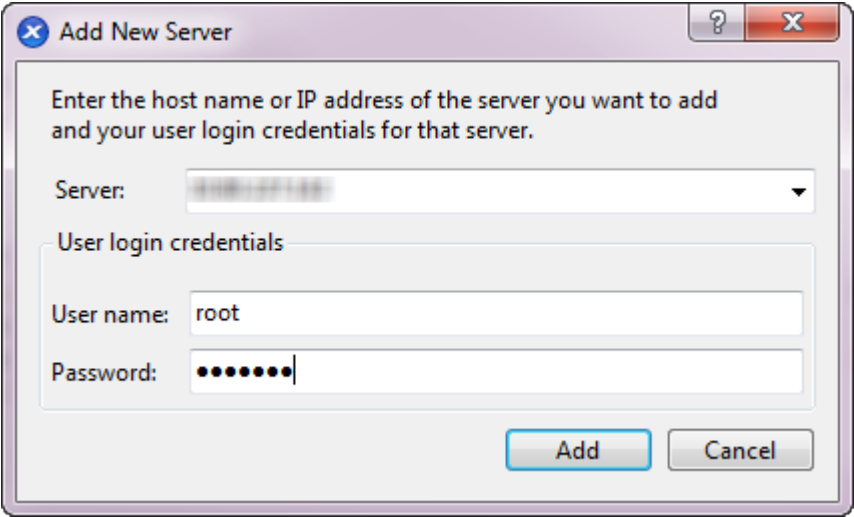
To connect XenCenter to the Citrix Hypervisor server:

1. Launch XenCenter.

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



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



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

Note:

The first time you add a server, the **Save and Restore Connection State** dialog box appears. This

dialog box enables you to set your preferences for storing your server connection information and automatically restoring server connections.

License Citrix Hypervisor

You can use Citrix Hypervisor without a license (Free Edition). However, this edition provides a restricted set of features.

If you have a Citrix Hypervisor license, apply it now.

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

Create a pool of Citrix Hypervisor servers

A resource pool is composed of multiple Citrix Hypervisor server installations, bound together as a single managed entity.

Resource pools enable you to view multiple servers and their connected shared storage as a single unified resource. You can flexibly deploy of VMs across the resource pool based on resource needs and business priorities. A pool can contain up to 64 servers running the same version of Citrix Hypervisor software, at the same patch level, and with broadly compatible hardware.

One server in the pool is designated as the *pool master*. The pool master 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 master if necessary. The pool master is the first server listed for the pool in the XenCenter Resources pane. You can find the pool master's IP address by selecting the pool master 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 servers. The VMs are moved while running and with minimal downtime. If an individual Citrix Hypervisor server suffers a hardware failure, you can restart the failed VMs on another server in the same pool.

If the high availability feature is enabled, protected VMs are *automatically* moved if a server fails. On an HA-enabled pool, a new pool master is automatically nominated if the master 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 servers
- Set up a network for the pool
- Bond NICs
- Set up shared storage for the pool

While Citrix Hypervisor 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 Citrix Hypervisor server, with similar processor type.
Connect this server to your XenCenter application.
- A storage repository for IP-based storage

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

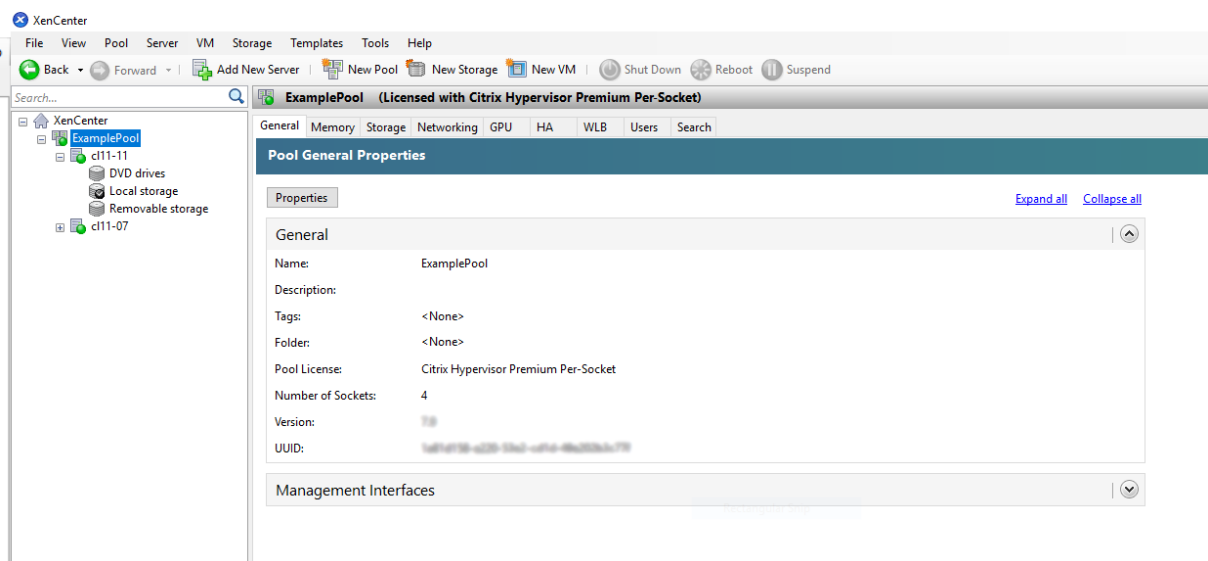
Create a pool

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 master by selecting a server from the **Master** list.
4. Select the second server 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 Citrix Hypervisor, you create a network connection, typically on the first NIC in the pool where you specified an IP address (during Citrix Hypervisor 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 Citrix Hypervisor to connect each NIC to one physical network and numerous VLANs.

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

Note:

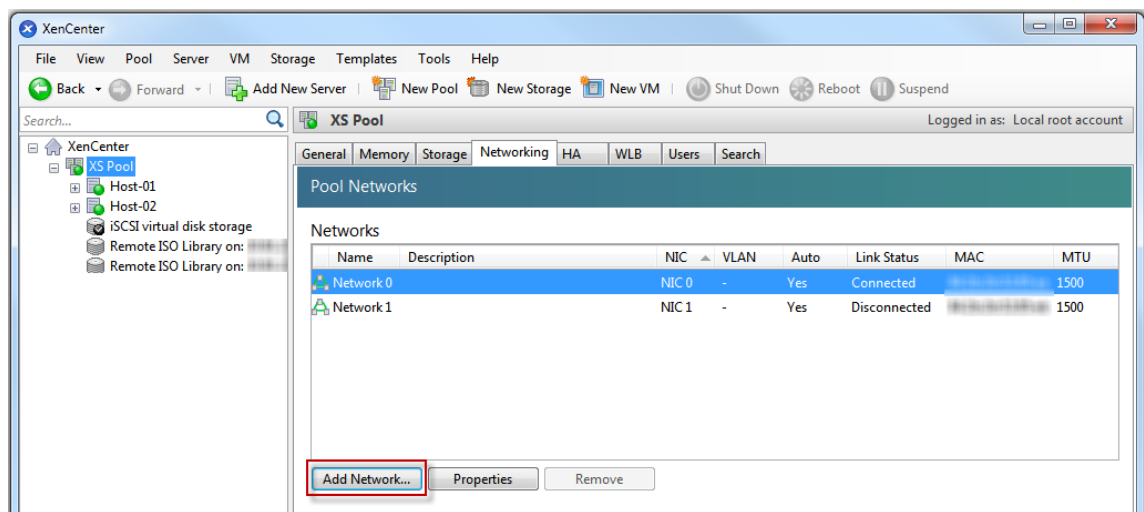
If the NICs were not plugged in to the NICs on the server when you installed Citrix Hypervisor:

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

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

To add a network to Citrix Hypervisor:

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 Citrix Hypervisor 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 server 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](#).

Citrix Hypervisor 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 options are not visible or available unless you configure the vSwitch as the network stack. Likewise, your switches must support the IEEE 802.3ad standard. The switch must contain a separate LAG group configured for each LACP bond on the server. 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 server 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 when vSwitch is the network stack. However, you can only bond two NICs when Linux bridge is the network stack.
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 server'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 server. 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.

Note:

LACP bonding is only available for the vSwitch, whereas active-active and active-passive bonding modes are available for both the vSwitch and Linux bridge. For information about networking stacks, see [Networking](#).

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 master when the new bond is created. A server with its management interface on a bond is not permitted to join a pool. Before the server 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 servers in a pool to a remote storage array, create a Citrix Hypervisor 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 Citrix Hypervisor. 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 Citrix Hypervisor 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 servers: 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 server. 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 Citrix Hypervisor
- 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 server). 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:

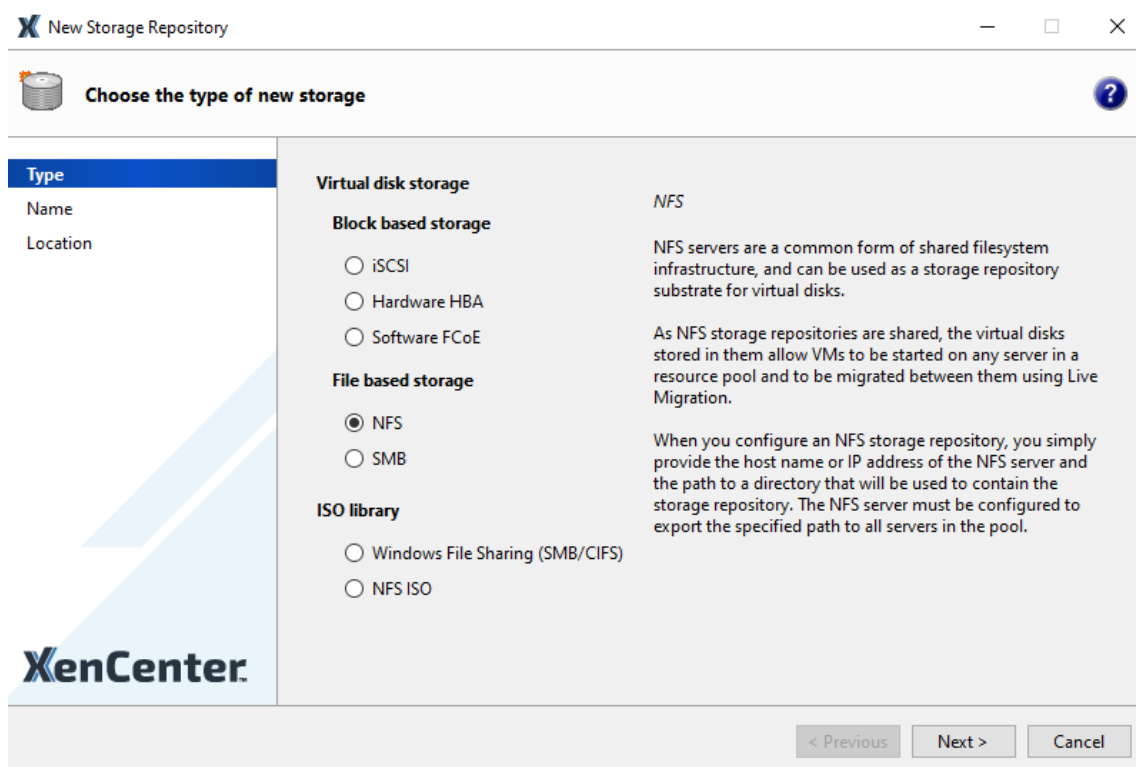
Whenever you assign a NIC an IP address, it must be on a different subnet than any other NICs with IP addresses in the pool. This includes the primary management interface.

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 server must be configured to export the specified path to all Citrix Hypervisor servers 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 Citrix Hypervisor server 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 servers 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.

Citrix Hypervisor also provides an easy way to convert batches of virtual machines from VMware. For more information, see [XenServer 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 Citrix Hypervisor: a single Citrix Hypervisor server with local storage (after you connect XenCenter to the Citrix Hypervisor server, storage is automatically configured on the local disk of the server).

This section also demonstrates how to use live migration to live migrate VMs between servers 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 VM
- Install XenServer VM Tools for Windows (formerly Citrix VM Tools)
- Migrate a running VM between servers 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 Citrix Hypervisor pool you set up
- XenCenter
- Installation files for Windows
- Installation files for XenServer VM Tools for Windows

Create a Windows 10 (32-bit) VM

Note:

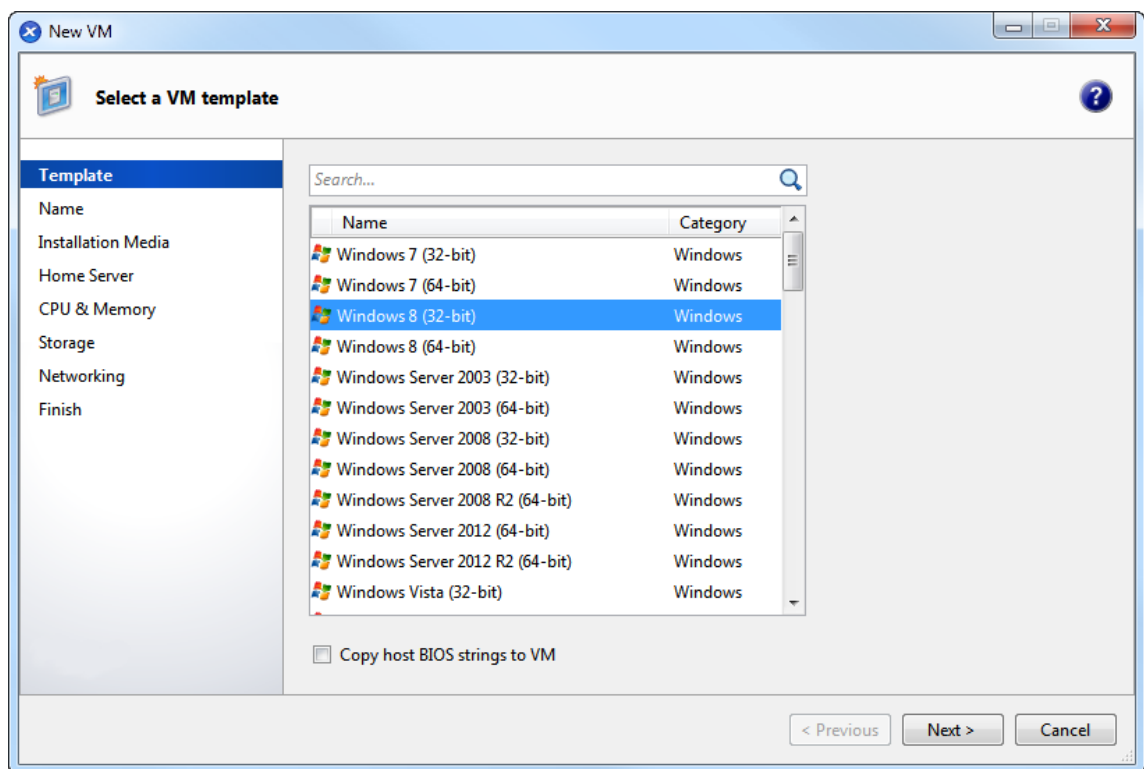
The following procedure provides an example of creating Windows 10 (32-bit) VM. The default values may vary depending on the operating system that you choose.

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 Citrix Hypervisor 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 Citrix Hypervisor server, and choose **Next** to proceed.

Citrix Hypervisor 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 server. Choose **Next** to proceed.

6. Allocate processor and memory resources.

For a Windows 10 VM, the default is 1 virtual CPU, 1 socket with 1 core per socket and 4 GB of RAM. You may choose to modify the defaults if necessary. Click **Next** to continue.

Note:

Each OS has different configuration requirements which are reflected in the templates.

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 Citrix Hypervisor Premium Edition customers, or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.

8. Configure storage for the new VM.

Click **Next** to select the default allocation (24 GB) 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 Citrix Hypervisor servers, 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:

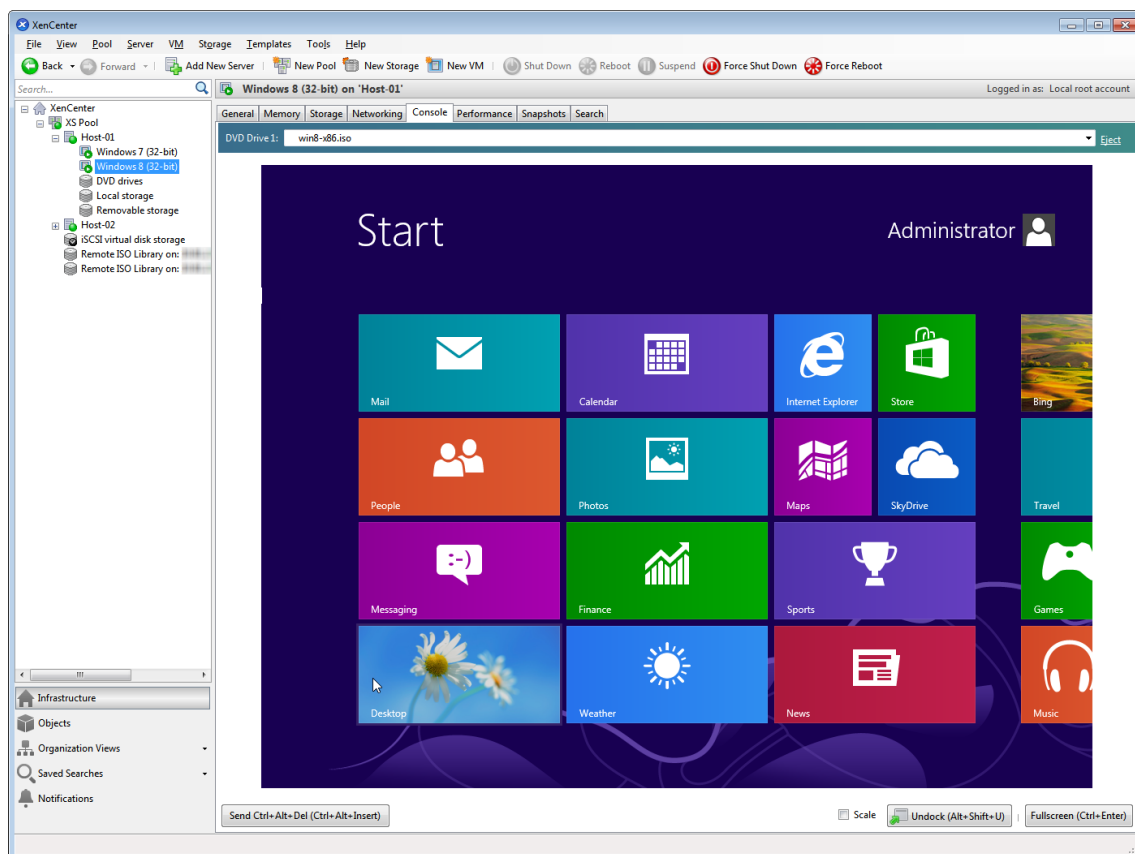
- a) Change the physical network, MAC address, or Quality of Service (QoS) priority of the virtual disk by clicking **Edit**.

- b) Add a new virtual network interface by selecting **Add**.

Citrix Hypervisor uses the virtual network interface to connect to the physical network on the server. 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 server 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 [Citrix Hypervisor downloads page](#).
A 32-bit and a 64-bit version of the tools are available.
2. Run the `managementagent.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 Citrix Hypervisor Premium Edition customers, or those customers who have access to Citrix Hypervisor through Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement:

- 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 servers in a pool

Using live migration, you can move a running VM from one server 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.

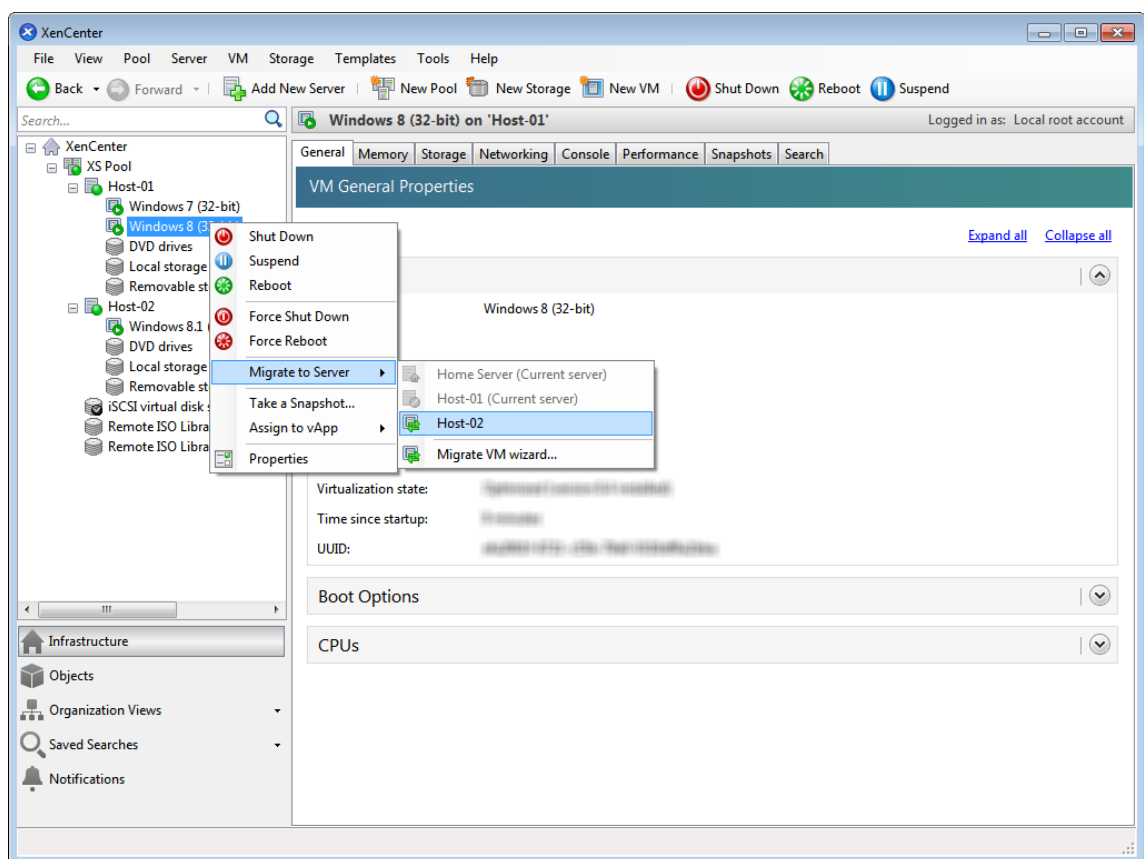
Note:

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 server.

Tip:

You can also drag the VM onto the target server.



3. The migrated VM displays under the new server 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:

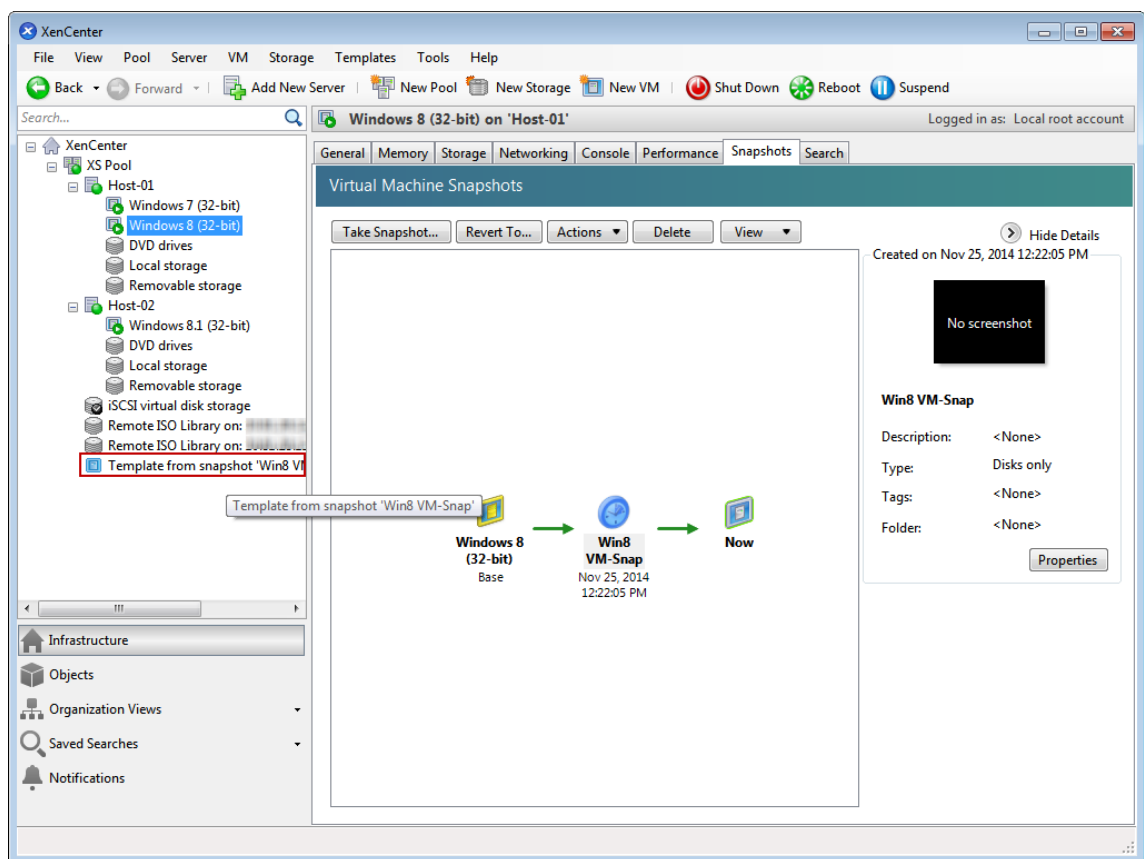
When you create a template from an existing VM, the new template replaces the original VM. The VM no longer exists.

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.



4. From the **Actions** list, choose **New Template from Snapshot**.
5. 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.

Technical overview

March 14, 2024

Citrix Hypervisor is an industry leading platform for cost-effective desktop, server, and cloud virtualization infrastructures. Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor hosts
- Assuring availability of VMs by using high availability to configure policies that restart VMs on another server in case one fails
- Increasing portability of VM images, as one VM image works on a range of deployment infrastructures

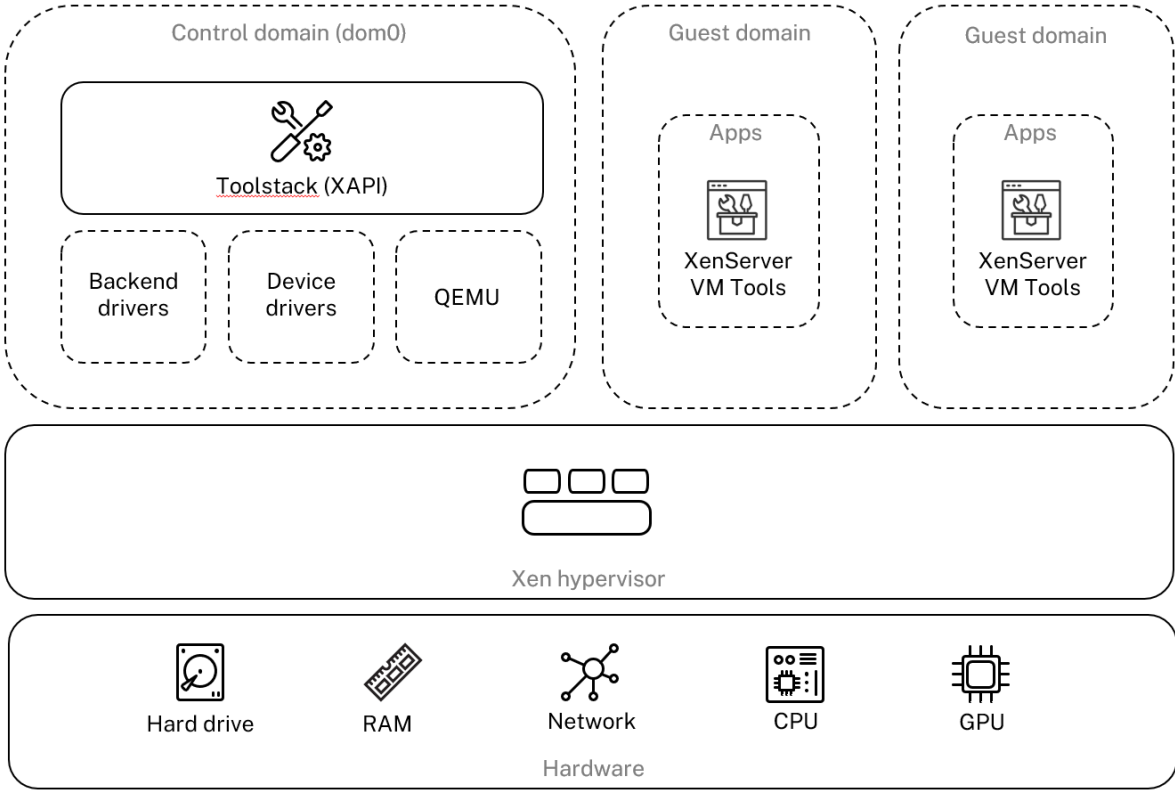
Virtualization and hypervisor

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 Citrix Hypervisor works. See the following illustration for the key components of Citrix Hypervisor:



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 Citrix Hypervisor host system requirements, see System requirements.

For a complete list of Citrix Hypervisor 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.

Citrix Hypervisor is based on the Xen Project hypervisor, with extra features and supports provided by Citrix. Citrix Hypervisor 8.2 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 Citrix Hypervisor management toolstack known as XAPI. This Linux VM is based on a CentOS 7.5 distribution. Besides providing Citrix Hypervisor 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.

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 Citrix Hypervisor 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 <https://developer.cloud.com/citrixworkspace/citrix-hypervisor/docs/overview>.

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. Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor. For a list of supported distributions using PV on HVM, see [Guest Operating System Support](#).

XenServer VM Tools XenServer VM Tools provide high performance I/O services without the overhead of traditional device emulation.

- XenServer VM Tools for Windows (formerly Citrix VM Tools) 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 Citrix Hypervisor 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.

- Citrix 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:

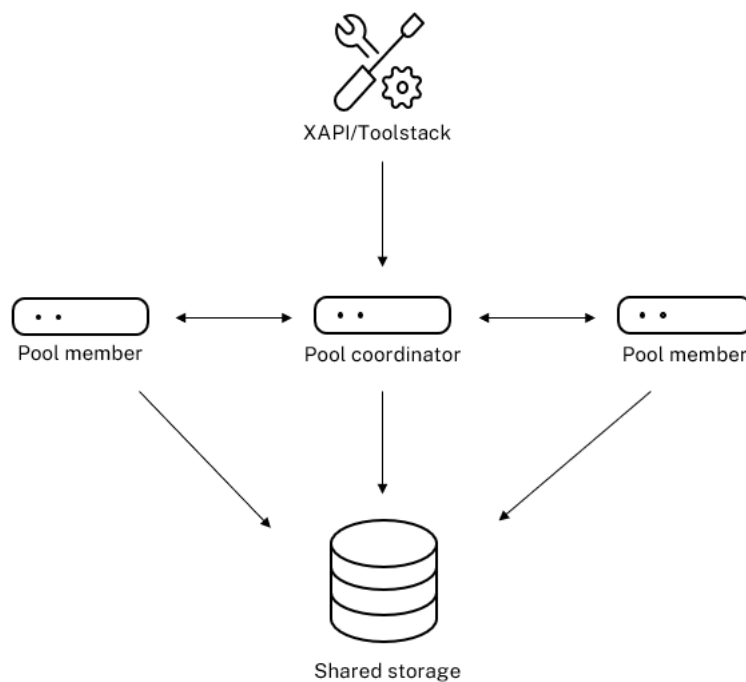
You cannot use the Dynamic Memory Control (DMC) feature on Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, Rocky Linux 8, Rocky Linux 9, or CentOS Stream 9 VMs as these operating systems do not support memory ballooning with the Xen hypervisor.

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

Key concepts

Resource pool

Citrix Hypervisor allows you to manage multiple servers 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 Citrix Hypervisor hosts. They also allow all servers to share a common framework for network and storage. A pool can contain up to 64 servers running the same version of Citrix Hypervisor software, at the same patch level, and with broadly compatible hardware. For more information, see [Hosts and resource pools](#).



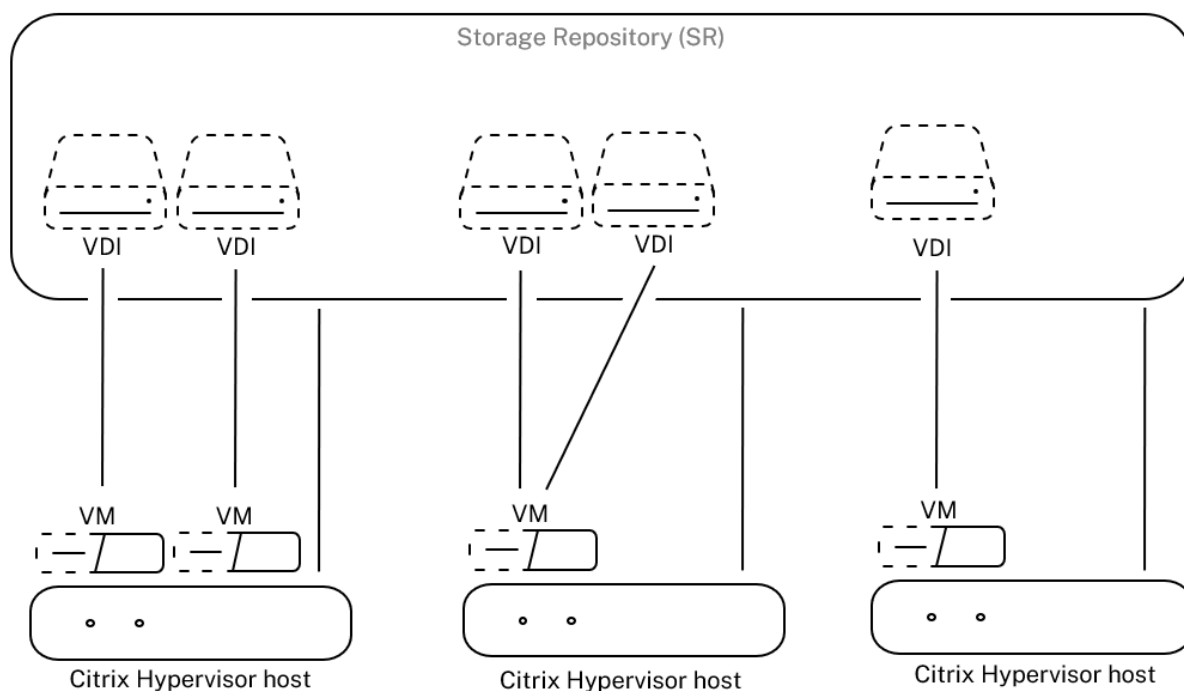
Resource pool overview

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

Storage repository

Citrix Hypervisor 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

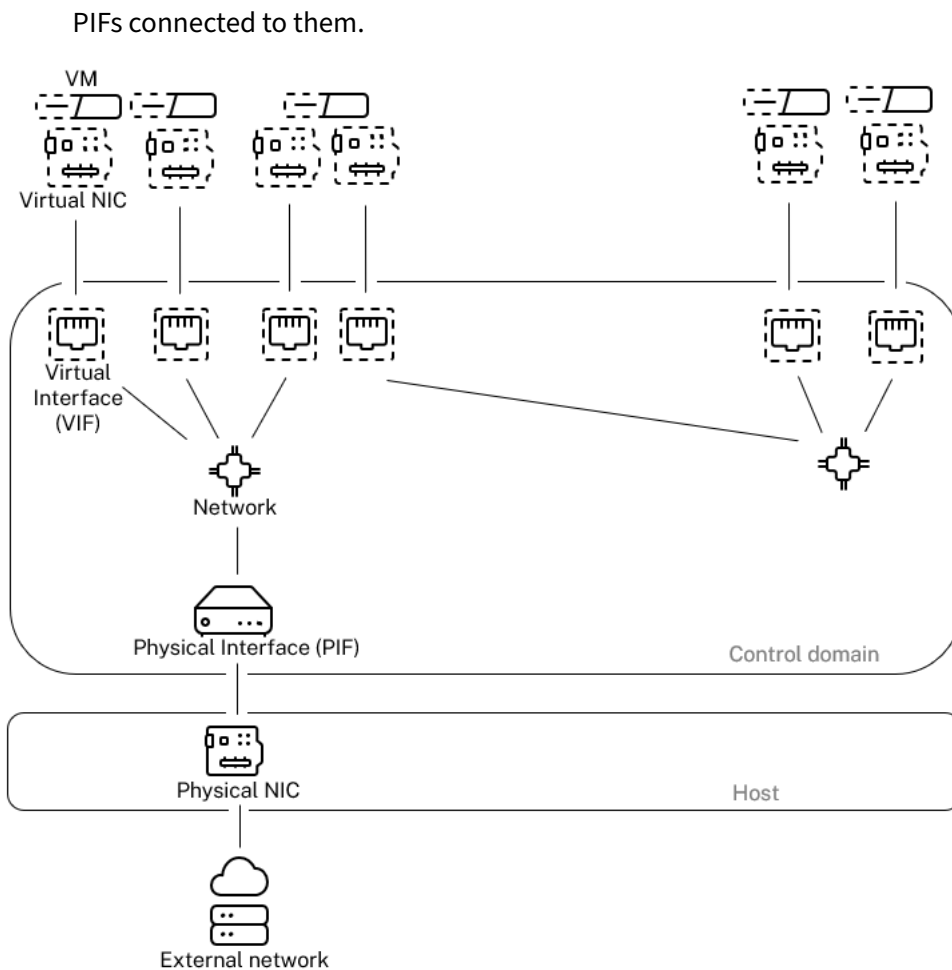
Each Citrix Hypervisor host can use multiple SRs and different SR types simultaneously. These SRs can be shared between hosts or dedicated to particular hosts. Shared storage is pooled between multiple hosts within a defined resource pool. A shared SR must be network-accessible to each host in the pool. 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 in dom0 and represents a physical NIC on a host. PIF objects have a name and description, a UUID, the parameters of the NIC that they represent, and the network and server they are connected to.
- A **VIF**, which is a software object used within in 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



Networking overview

Citrix Hypervisor 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 Citrix Hypervisor, see [Networking](#).

Related add-ons and applications

While Xen Hypervisor works at the core level, there are Citrix Hypervisor specific add-ons related hypervisor-agnostic applications and services available to make the virtualization experience complete.

- **XenCenter**

A windows GUI client for VM management, implemented based on the management API. XenCenter provides a rich user experience to manage multiple Citrix Hypervisor 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 servers for their workload in a resource pool. For more information, see [Workload balancing \(/en-us/citrix-hypervisor/wlb.html\)](/en-us/citrix-hypervisor/wlb.html).

- **Citrix Licensing Server**

A Linux based appliance that XenCenter contacts to request a license for the specified server.

- **XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager)**

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

- **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 Citrix Hypervisor in a multi-host pool configuration. For more information, see [Citrix Virtual Apps and Desktops](#).

Technical FAQs

July 22, 2024

Hardware

What are the minimum system requirements for running Citrix Hypervisor?

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

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

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 BIOS.

What systems are certified to run Citrix Hypervisor?

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

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

Yes. Citrix Hypervisor 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 Citrix Hypervisor run on a notebook or desktop-class systems?

Citrix Hypervisor runs on many notebook or desktop-class systems that conform to the minimum CPU requirements. However, Citrix 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 Citrix Hypervisor be installed on SD or USB cards?

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

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

Product limits

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

Resource sharing

How are processing resources split between VMs?

Citrix Hypervisor 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 Citrix Hypervisor choose which physical processors it allocates to the VM?

Citrix Hypervisor doesn't statically allocate physical processors to any specific VM. Instead, Citrix Hypervisor 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?

Citrix Hypervisor 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?

Citrix Hypervisor uses a fair share resource split for network I/O resources between the VMs. You can also control bandwidth-throttling limits per VM by using the Open vSwitch.

Guest operating systems

Can Citrix Hypervisor run 32-bit operating systems as guests?

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

Can Citrix Hypervisor 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 Citrix Hypervisor?

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

Which versions of Linux can run as guests on Citrix Hypervisor?

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?

Citrix only supports 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, Citrix can only support the operating systems listed in [Supported guest operating systems](#). We might ask you to switch to a supported OS before we can investigate any issues.

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

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

What are the XenServer VM Tools?

The XenServer VM Tools (formerly Citrix 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 Citrix Hypervisor 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 Citrix Hypervisor.

Can I run Docker containers on my Windows VMs?

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

Does Citrix Hypervisor provide additional features for working with Docker?

No.

In previous releases of Citrix Hypervisor and XenServer, a Container Management supplemental pack was available that enabled you to manage your Docker containers through XenCenter. This feature has been removed.

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 Citrix Hypervisor 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 Citrix Hypervisor.

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 Citrix Hypervisor hosts?

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

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

Yes. XenCenter is backward compatible with multiple host systems running different versions of Citrix Hypervisor that are currently supported.

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 Citrix Hypervisor include a CLI?

Yes. All editions of Citrix Hypervisor include a full command line interface (CLI) –known as [xe](#).

Can I access the Citrix Hypervisor 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 Citrix Hypervisor CLI from a remote system?

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

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

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

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

Yes. You can restrict user access on the Citrix Hypervisor CLI.

VMs

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

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

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

You can also convert VMs in batches using the XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager). Third-party tools are also available.

For more information, see [XenServer 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 using technology such as DRAC
- 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 Citrix Hypervisor 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 Citrix Hypervisor and moved to another?

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

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

No.

Does Citrix Hypervisor provide disk snapshot capabilities for VMs?

Yes. Citrix Hypervisor 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 Citrix Hypervisor?

Citrix Hypervisor supports local storage such as SATA, SAS, and NVMe.

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

Citrix Hypervisor 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 Citrix Hypervisor support software-based iSCSI?

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

What version of NFS is required for remote storage use?

Citrix Hypervisor requires NFSv3 or NFSv4 over TCP for remote storage use. Citrix Hypervisor 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 Citrix recommends 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 Citrix Hypervisor host system from an iSCSI, Fibre Channel or FCoE SAN?

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

Can I boot a Citrix Hypervisor host using UEFI?

Yes. Citrix Hypervisor supports booting from BIOS and UEFI. However, UEFI Secure Boot is not supported for the Citrix Hypervisor host.

For more information, see [Network boot installations](#)

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

Yes. Citrix recommends using multipath for resilient storage connections.

Does Citrix Hypervisor support a software-based RAID implementation?

No. Citrix Hypervisor doesn't support software RAID.

Does Citrix Hypervisor support HostRAID or FakeRAID solutions?

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

Does Citrix Hypervisor 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 Citrix Hypervisor support Distributed Replicated Block Device (DRBD) storage?

No. Citrix Hypervisor doesn't support DRBD.

Does Citrix Hypervisor support ATA over Ethernet?

No. Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor support IPv6?

VMs hosted on Citrix Hypervisor can use any combination of IPv4 and IPv6 configured addresses.

However, Citrix Hypervisor 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 Citrix Hypervisor host to use.

Does Citrix Hypervisor support VLANs on a physical network interface?

Yes. Citrix Hypervisor supports assigning VM networks to specified VLANs.

Do Citrix Hypervisor virtual networks pass all network traffic to all VMs?

No. Citrix Hypervisor uses Open vSwitch (OVS), which acts as a Layer 2 switch. A VM only sees traffic for that VM. Also, the multitenancy support in Citrix Hypervisor enables increased levels of isolation and security.

Do the virtual network interfaces and networks support promiscuous mode?

Yes. If you are using the Linux bridge as the network stack, your virtual network interfaces can be configured for promiscuous mode. This mode enables you to see all traffic on a virtual switch. For more information about promiscuous mode configuration, see the following Knowledge Center articles:

- [CTX116493 - How to Enable Promiscuous Mode on a Physical Network Card](#)
- [CTX121729 - How to Configure a Promiscuous Virtual Machine in XenServer](#)

When you enable promiscuous mode on a virtual network interface, for a VM to make use of this configuration, you must also enable promiscuous mode within your VM.

Does Citrix Hypervisor support bonding or teaming of physical network interfaces?

Yes. Citrix Hypervisor 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 Citrix Hypervisor?

Three components contribute to the memory footprint of a Citrix Hypervisor host.

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

The amount of memory required to run dom0 is adjusted automatically. By default, Citrix Hypervisor 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 Citrix Hypervisor Crash Kernel. The amount of memory used by the hypervisor is larger for hosts with more memory.

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

Does Citrix Hypervisor optimize VM memory usage?

Yes. Citrix Hypervisor uses Dynamic Memory Control (DMC) to adjust automatically 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 Citrix Hypervisor 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 Citrix Hypervisor require a dedicated host to manage a resource pool?

No. A single host in the pool must be specified as the Pool Master. The Pool Master controls all administrative activities required on the pool. This design means that there is no external single point

of failure. If the Pool Master fails, other hosts in the pool continue to operate, and the resident VMs continue to run as normal. If the Pool Master cannot come back online, Citrix Hypervisor promotes one of the other hosts in the pool to master 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 master fails, this data enables any host in the resource pool to become the new pool master.

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 master 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 Citrix Hypervisor resource pool?

Yes. Citrix recommends 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 Citrix Hypervisor offer high availability features?

Yes. If high availability is enabled, Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor management tools collect performance information?

Yes. Citrix Hypervisor 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 Citrix Hypervisor performance metrics gathered?

Data for the Citrix Hypervisor 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 Citrix Hypervisor host. You can customize the metrics that are displayed.

Does XenCenter store and display historic performance metrics?

Yes. Citrix Hypervisor 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 Citrix Hypervisor install on top of systems that are already running an existing operating system?

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

Can I upgrade an existing Citrix Hypervisor installation to a newer version?

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

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

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

You must create a fresh installation of Citrix Hypervisor 8.2 Cumulative Update 1. Any other upgrade path for these out-of-support versions is not supported.

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

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

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

Yes. You can install Citrix Hypervisor on the host system using PXE. You can also automatically install Citrix Hypervisor 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 Citrix Hypervisor get its device driver support?

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

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

Licensing

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

Technical Support

Does Citrix provide direct technical support for Citrix Hypervisor?

Yes. For more information, visit [Citrix Support and Services](#).

Can I get technical support for Citrix Hypervisor and other Citrix products on a single support contract?

Yes. Citrix provides Technical Support contracts that allow you to open support incidents on Citrix Hypervisor, in addition to other Citrix products.

For more information, visit [Citrix Support and Services](#).

Do I have to buy a Citrix technical support contract at the same time as I buy Citrix Hypervisor?

No. You can buy a technical support contract from Citrix either at product point-of-sale or at another time.

Are there alternative channels for getting technical support for Citrix Hypervisor?

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

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

No. Citrix 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 Citrix Hypervisor release reaches EOL?

We support Citrix Hypervisor in accordance with the dates specified in the Citrix Product Matrix (<https://www.citrix.com/support/product-lifecycle/product-matrix.html>). 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 Citrix if I'm experiencing a non-technical issue?

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

Licensing

July 15, 2024

Citrix Hypervisor 8.2 Cumulative Update 1 is available in the following editions:

- Standard Edition
- Premium Edition

The **Standard Edition** is our entry-level commercial offering. It has a range of features for customers who want a robust and high performing virtualization platform, but don't require the premium features of Premium Edition. Meanwhile, they still want to benefit from the assurance of comprehensive Citrix Support and Maintenance.

The **Premium Edition** is our premium offering, optimized for desktop, server, and cloud workloads. In addition to the features available in the Standard Edition, the Premium Edition offers the following features:

- Automated Windows VM driver updates
- Automatic updating of the Management Agent
- Support for SMB storage
- Direct Inspect APIs
- Dynamic Workload Balancing
- GPU virtualization with NVIDIA vGPU, AMD MxGPU, and Intel GVT-g
- VMware vSphere to Citrix Hypervisor conversion utilities
- Export pool resource data
- In-memory read caching
- PVS-Accelerator
- Citrix Hypervisor live patching
- Enablement for Citrix Virtual Desktops tablet mode
- Changed block tracking
- IGMP snooping
- USB pass-through
- SR-IOV network support
- Thin provisioning for shared block storage devices

Note:

Automated Updates were previously restricted to Citrix Hypervisor Premium Edition customers or Citrix Virtual Apps and Desktops customers. However, in pools with hotfix XS82ECU1053 applied, this feature is available to all users.

Customers who have purchased Citrix Virtual Apps or Citrix Virtual Desktops or who use Citrix DaaS (formerly Citrix Virtual Apps and Desktops service) with on-premises desktops and apps have an entitlement to Citrix Hypervisor, which includes all the features listed for Premium Edition.

For more information, see the [Citrix Hypervisor Feature Matrix](#).

Citrix Hypervisor uses the same licensing process as other Citrix products, and as such requires a valid license to be installed on a License Server. You can download the License Server for Windows from [Citrix Licensing](#). Citrix Hypervisor (other than through the Citrix Virtual Apps and Desktops licenses) is licensed on a *per-socket* basis. Allocation of licenses is managed centrally and enforced by a stand-alone Citrix License Server in the environment. After applying a per socket license, Citrix Hypervisor displays as *Citrix Hypervisor Per-Socket Edition*.

An unlicensed or Express Edition is not available for Citrix Hypervisor 8.2 Cumulative Update 1.

Licensing steps overview

You need the following items to license Citrix Hypervisor Premium Edition or Standard Edition:

- A Citrix Hypervisor License
- A Citrix License Server
- A Citrix Hypervisor server
- XenCenter

Note:

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

The following steps provide an overview of the process:

1. Install Citrix License Server or import the Citrix License Server virtual appliance into a Citrix Hypervisor server.
2. Download a license file that is tied to the case-sensitive host name of your License Server.
3. Add the license file to the Citrix License Server.
4. Using XenCenter, enter the License Server details and apply them to hosts in your resource pool.

This process is covered in detail in the Citrix Licensing documentation. For more information, see [Licensing guide for Citrix Hypervisor](#).

Licensing Citrix Hypervisor

Q: Where can I get a Citrix Hypervisor license?

A: You can speak to a Citrix representative about getting a license at [the Citrix website](#).

Q: How do I apply a Citrix Hypervisor license?

A: Citrix Hypervisor requires a License Server. After licensing Citrix Hypervisor, you are provided with a .LIC license access code. Install this license access code on either:

- A Windows server running the Citrix License Server software
- The Citrix License Server virtual appliance

When you assign a license to a Citrix Hypervisor server, Citrix Hypervisor contacts the specified Citrix License Server and requests a license for the specified servers. If successful, a license is checked out and the License Manager displays information about the license the hosts are licensed under.

Q: How many licenses do I need to license my resource pool?

A: Citrix Hypervisor is licensed on a *per-CPU socket* basis. For a pool to be considered licensed, all Citrix Hypervisor servers in the pool must be licensed. Citrix Hypervisor 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: 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 30 days at the licensing level that was applied previously.

Q: I am upgrading to Citrix Hypervisor 8.2 from a previous Citrix Hypervisor version with a per socket license. Do I have to do anything?

A: No. You can upgrade your hosts to Citrix Hypervisor 8.2 Premium Edition using the previously bought per socket licenses, provided Customer Success Services is valid at least until Dec 13, 2021.

If you have renewed your Customer Success Services after the original purchase, you might need to refresh the license file on the License Server to ensure it displays the Customer Success Services eligibility.

Q: I am moving from Citrix Hypervisor 8.2 Express Edition (unlicensed) to Citrix Hypervisor 8.2 Cumulative Update 1. Do I have to do anything?

A: Yes. You must apply an appropriate license to update your hosts to Citrix Hypervisor 8.2 Cumulative Update 1.

An Express Edition is not available for Citrix Hypervisor 8.2 Cumulative Update 1.

Q: What happens if my license expires?

We strongly recommend that you do not let your license expire.

If you let your license expire on Citrix Hypervisor 8.2 CU1, you lose access to important features, including but not limited to:

- All Premium Edition features
- The ability to add hosts to pools with three or more members
- Live migration and storage live migration
- Rolling pool upgrade
- GPU passthrough
- High availability and disaster recovery
- Dynamic memory control
- Active Directory integration
- Role-based access control

Any VMs that are already running continue until they are stopped. Stopped VMs cannot be started.

Because of restrictions on features, you might experience unexpected behaviors.

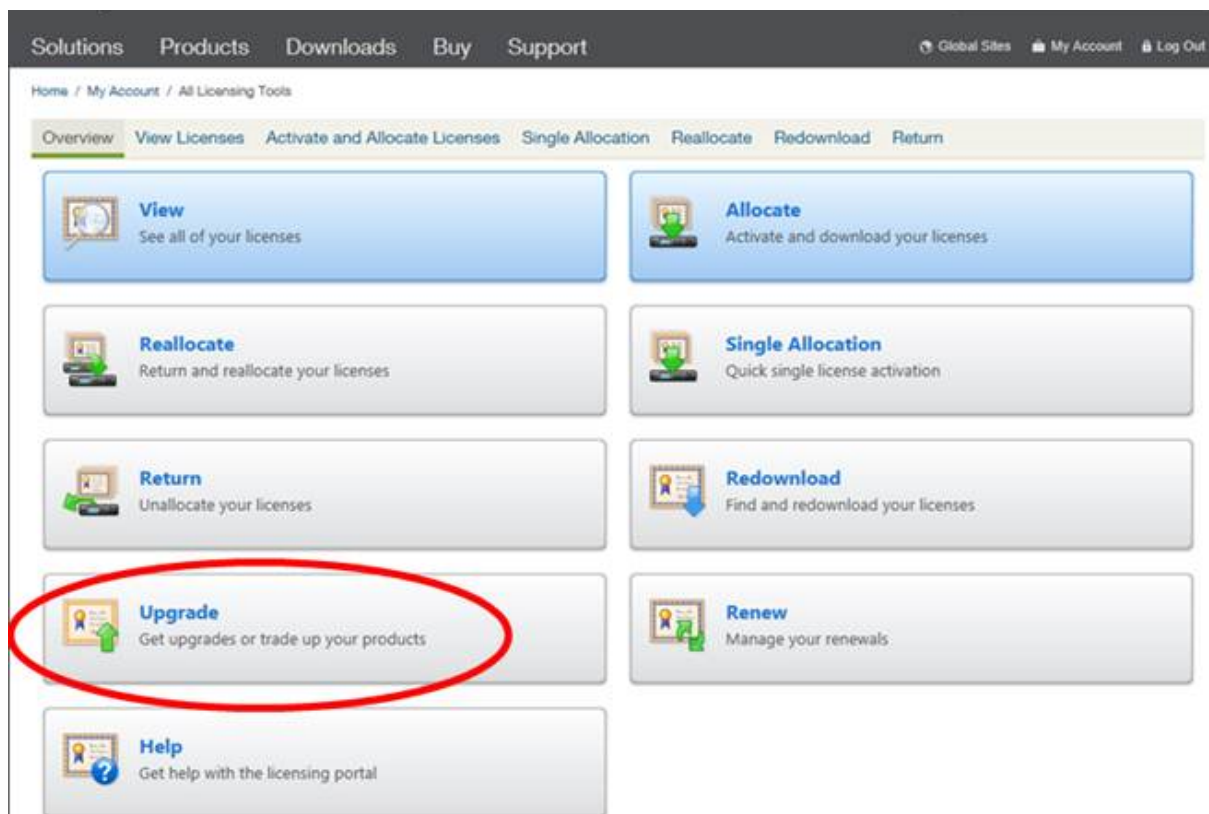
We do not provide support for issues caused by an expired license.

Q: Do I lose my virtual machine (VM) when my license expires?

A: No, you do not lose any VMs or their data.

Q: I am a Citrix Virtual Apps and Desktops or Citrix DaaS customer moving from an earlier version of Citrix Hypervisor to Citrix Hypervisor 8.2. Do I have to do anything?

A: No. Citrix Virtual Apps, Citrix Virtual Desktops, or Citrix DaaS customers can update to Citrix Hypervisor 8.2 seamlessly. Your existing installed Citrix Virtual Apps or Citrix Virtual Desktops license grants you entitlement to Citrix Hypervisor without requiring any other changes, provided your license is valid at least until Jun 25, 2020.

**Q: I am a Citrix Service Provider licensed for Citrix Virtual Apps and Desktops or Citrix DaaS. Can I use this license for Citrix Hypervisor when I upgrade to Citrix Hypervisor 8.2?**

A: Yes. Citrix Hypervisor 8.2 supports your license. With this license, you can use all the premium features provided by the Premium Edition of Citrix Hypervisor. To apply this license to your pools, first upgrade or update all hosts within the pool to run Citrix Hypervisor 8.2.

Q: I am a customer with a Citrix DaaS subscription. Am I entitled to use Citrix Hypervisor 8.2?

A: Yes. If you have a Citrix DaaS (formerly Citrix Virtual Apps and Desktops Service) subscription that enables the use of on-premises desktops and apps, you are entitled to use Citrix Hypervisor to host these desktops and apps.

Download a license through the licensing management tool. Install this license on your License Server to use an on-premises Citrix Hypervisor with your Citrix DaaS subscription.

With this license you can use all the same premium features as with an on-premises Citrix Virtual Apps and Desktops entitlement. To apply this license to your pools, first upgrade all hosts within the pool to run Citrix Hypervisor 8.2.

Q: What are the constraints on the use of the Citrix Hypervisor Premium Edition advanced virtualization management capabilities delivered as part of Citrix Virtual Apps and Desktops?

A: Every edition of Citrix Virtual Apps and Desktops has access to Citrix Hypervisor Premium Edition advanced virtualization management features. A complete list of all features enabled by a Citrix Virtual Apps or Citrix Virtual Desktops license can be found in the Citrix Hypervisor Feature Matrix.

Citrix Hypervisor entitlements permit virtualization of any infrastructure required to deliver Citrix Virtual Apps or Citrix Virtual Desktops feature components. These features must be accessed exclusively by Citrix Virtual Apps or Citrix Virtual Desktops licensed users or devices.

Extra infrastructure support servers, such as Microsoft domain controllers and SQL servers are also covered by this entitlement, providing they are deployed in the same Citrix Hypervisor resource pool as the Citrix Virtual Apps or Citrix Virtual Desktops infrastructure covered by this license, and providing those support servers are used to support the Citrix Virtual Apps or Citrix Virtual Desktops infrastructure only.

The Citrix Hypervisor entitlement in the Citrix Virtual Apps or Citrix Virtual Desktops license cannot be used for Citrix Hypervisor pools that don't host Citrix Virtual Apps or Citrix Virtual Desktops infrastructure or Virtual Delivery Agents (VDAs). You also cannot use this entitlement for hosting virtual machines not covered by the permissions above. Citrix Hypervisor must be purchased separately for these uses.

Citrix License Servers

Q: Which License Servers can I use with Citrix Hypervisor?

A: You can use the Citrix License Server software version 11.16 or later on a server running Microsoft Windows.

Previously, we supported the License Server virtual appliance. However, the License Server virtual appliance is now out of support and won't receive any further maintenance or security fixes. Customers using 11.16.6 or previous versions of License Server virtual appliance are advised to migrate to the latest version of License Server for Windows as soon as possible.

Q: How do I import my license onto the Citrix License Server?

A: For information on importing a license file, see the [Citrix License Server documentation](#).

Q: Can I run the License Server on my Citrix Hypervisor pool?

A: Yes. You can install the Citrix License Server software on a Windows VM.

Citrix Hypervisor operates with a ‘grace’ license until the License Server is able to boot. This behavior means, after you have licensed the Citrix Hypervisor servers 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 Citrix Hypervisor?

A: Yes.

Q: Can I install Licenses for other Citrix products on a Citrix License Server virtual appliance or on the Citrix License Server software installed on Windows?

A: Yes, you can license other Citrix products by using the Citrix License Server virtual appliance or through the Citrix License Server software installed on Windows. For more information, see [Licensing](#) on the [Citrix Product Documentation](#) website.

Licensing a Citrix Hypervisor 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 would like 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=enterprise-per-socket|desktop-plus|  
  desktop|standard-per-socket \  
2  
3 license-server-address=<licenseserveraddress> host-uuid=<  
  uuidofhost> \  
4  
5 license-server-port=<licenseserverport>
```

To license a pool, use the pool-apply-edition command. For example:

```
1 xe pool-apply-edition edition=enterprise-per-socket|desktop-plus|  
  desktop|standard-per-socket \  
2  
3 license-server-address=<licenseserveraddress> pool-uuid=<  
  uuidofpool> \  
4  
5 license-server-port=<licenseserverport>
```

Q: How can I discover the license status of my servers and pools?

A: XenCenter displays the license type of a server or pool.

To see the license type of a server or pool, select that server or pool in the tree view. XenCenter displays the license status in the title bar for that server or pool, after the server or pool name.

You can also go to the **General** tab of the server and find the license type in the **License Details** section.

To find the license type of a server by using the command line, run the following command in the console of a server in your pool:

```
1 xe host-license-view host\_uuid=<UUID> | grep sku\_marketing\_name
```

Other questions

Q: How do I get a license to evaluate Citrix Hypervisor?

A: You can request a demo of the Premium Edition features. For more information, see [Get started](#).

Q: Can I use Citrix Hypervisor 8.2 Cumulative Update 1 without a license?

A: No. Citrix Hypervisor 8.2 Cumulative Update 1 is only provided in licensed editions or for customers with a Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. An Express Edition is not available for Citrix Hypervisor 8.2 Cumulative Update 1.

You can request a demo of Citrix Hypervisor Premium Edition. For more information, see [Getting started with Citrix Hypervisor](#).

Q: Can I use Citrix Cloud to apply a license to Citrix Hypervisor?

No, Citrix Hypervisor does not support licensing hosted on Citrix Cloud. To license Citrix Hypervisor, you require a License Server. For more information, see [Licensing your hosts and pools](#).

More information

- For more information about the Citrix Hypervisor 8.2 release, see [Citrix Hypervisor 8.2 Release Notes](#).
- To access Citrix Hypervisor 8.2 product documentation, see [Citrix Hypervisor 8.2 Product Documentation](#).
- For an overview of the Citrix Hypervisor product, see [Technical Overview](#).
- Raise any non-technical issues with Citrix Hypervisor including Customer Success Services program support, licensing, administrative support, and order confirmation through [Citrix Customer Service](#).

Install

July 8, 2024

This section contains procedures to guide you through the installation, configuration, and initial operation of Citrix Hypervisor. It also contains information about troubleshooting problems that might occur during installation and points you to extra resources.

This information is primarily aimed at system administrators who want to set up Citrix Hypervisor servers on physical servers.

Citrix Hypervisor installs directly on bare-metal hardware avoiding the complexity, overhead, and performance bottlenecks of an underlying operating system. It uses the device drivers available from the Linux kernel. As a result, Citrix Hypervisor 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 Citrix Hypervisor server must be installed on a dedicated 64-bit x86 server. Do not install any other operating system in a dual-boot configuration with the Citrix Hypervisor server. This configuration is not supported.

Before you start

Before installing Citrix Hypervisor 8.2 Cumulative Update 1, consider the following factors:

- What is the appropriate installation method?
- What are the system requirements?

Installation methods

Citrix Hypervisor 8.2 Cumulative Update 1 can be installed as a fresh installation.

There is no supported direct upgrade path from out-of-support versions of XenServer to Citrix Hypervisor 8.2 Cumulative Update 1. Instead, perform a fresh installation.

Fresh installation If you are creating a fresh installation of Citrix Hypervisor 8.2 Cumulative Update 1:

- Use the **Citrix Hypervisor 8.2 Cumulative Update 1 Base Installation ISO** file. You can download this file from the [Citrix download site](#)
- Review the information in [System Requirements](#), [Licensing Citrix Hypervisor](#), and [Installing Citrix Hypervisor and XenCenter](#) before installing Citrix Hypervisor.

Supplemental packs

You can install any required supplemental pack after installing Citrix Hypervisor. Download the supplemental pack to a known location on your computer and install the supplemental pack in the same way as an update.

For more information, see [Supplemental Packs and the DDK Guide](#).

Install the Citrix Hypervisor server

Tip:

Throughout the installation, quickly advance to the next screen by pressing F12. Use Tab to move between elements and Space or Enter to select. Press F1 for general help.

To install or upgrade the Citrix Hypervisor server:

1. Back up data you want to preserve. Installing Citrix Hypervisor overwrites data on any hard drives that you select to use for the installation.
2. Boot the computer from the installation media or by using network boot:
 - To install Citrix Hypervisor server 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.
 - 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) In the BIOS, 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 Citrix Hypervisor server from a CD:
 - a) Burn the Citrix Hypervisor installation ISO file to a CD.
 - b) Insert the bootable CD into the CD/DVD drive on the target system.
 - c) Restart the system.
 - d) In the BIOS, change the settings to boot the system from the CD.
(If necessary, see your hardware vendor documentation for information on changing the boot order)
 - To install Citrix Hypervisor server from virtual media:
 - a) Go to the virtual console of your system.
 - b) Insert the Citrix Hypervisor installation ISO file as virtual media.
 - c) Restart the system.

d) In the BIOS, 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 set up a network-accessible TFTP server to boot:

For details about setting up a TFTP server to boot the installer using network, see [Network Boot Installation](#).

- To install Citrix Hypervisor to a remote disk on a SAN to enable boot from SAN:

For details, see [Boot From SAN](#).

3. Following the initial boot messages and the Welcome to Citrix Hypervisor 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. The Welcome to Citrix Hypervisor Setup screen is displayed.

Citrix Hypervisor ships with a broad driver set that supports most modern server hardware configurations. However, you might need to apply driver disks (a type of supplemental pack) in order to be able to perform the Citrix Hypervisor installation. If you have been provided with any additional essential device drivers, press F9. The installer steps you through loading the necessary drivers.

Warning:

You cannot install other types of supplemental packs at this point in the installation process. You can install them along with additional driver disks near the end of the installation process.

After you have loaded all of the required drivers, select **OK** to proceed.

Citrix Hypervisor enables customers to configure the Citrix Hypervisor installation to boot from FCoE. Press F10 and follow the instructions displayed on the screen to set up FCoE.

Notes:

Before enabling your Citrix Hypervisor server to boot from FCoE, manually complete the configuration required to expose a LUN to the host. This manual configuration includes configuring the storage fabric and allocating LUNs to the public world wide name (PWWN) of your SAN. After you complete this configuration, the available LUN is mounted to the CNA of the host as a SCSI device. The SCSI device can then be used to access the LUN as

if it were a locally attached SCSI device. For information about configuring the physical switch and the array to support FCoE, see the documentation provided by the vendor.

When you configure the FCoE fabric, do not use VLAN 0. The Citrix Hypervisor server cannot find traffic that is on VLAN 0.

Occasionally, booting a Citrix Hypervisor server from FCoE SAN using software FCoE stack can cause the host to stop responding. This issue is caused by a temporary link disruption in the host initialization phase. If the host fails to respond for a long time, you can restart the host to work around this issue.

5. The Citrix Hypervisor EULA is displayed. Use the Page Up and Page Down keys to scroll through and read the agreement. Select **Accept EULA** to proceed.
6. Select the appropriate action. You might see any of the following options:
 - *Perform clean installation*
 - *Upgrade:* If the installer detects a previously installed version of Citrix Hypervisor or XenServer, it offers the option to upgrade. For information about upgrading your Citrix Hypervisor server, see [Upgrading from an existing version](#).
 - *Restore:* If the installer detects a previously created backup installation, it offers the option to restore Citrix Hypervisor from the backup.

Make your selection, and choose **OK** to proceed.

7. If you have multiple local hard disks, choose a Primary Disk for the installation. Select **OK**.
8. Choose which disks you want to use for virtual machine storage. Information about a specific disk can be viewed by pressing **F5**.

If you want to use thin provisioning to optimize the use of available storage, select **Enable thin provisioning**. This option selects the local SR of the host to be the one to be used for the local caching of VM VDI. Citrix Virtual Desktops and DaaS users are recommended to select this option for local caching to work properly. For more information, see [Storage](#).

Choose **OK**.

9. Select your installation media source.
 - To install from a USB or CD, choose **Local media**.
 - To install by using network, select **HTTP** or **FTP** or **NFS**.

Note:

If you are using IIS to host the installation media, ensure that double escaping is enabled on IIS before extracting the installation ISO on it.

Choose **OK** to proceed.

If you select **HTTP** or **FTP** or **NFS**, set up networking so that the installer can connect to the Citrix Hypervisor installation media files:

- a) If the computer has multiple NICs, select one of them to be used to access the Citrix Hypervisor 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.

Select **OK** to proceed.

10. Indicate if you want to verify the integrity of the installation media. If you select **Verify installation source**, the SHA256 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.
11. Set and confirm a root password, which XenCenter uses to connect to the Citrix Hypervisor server. You also use this password (with user name “root”) to log into **xsconsole**, the system configuration console.

Note:

Citrix Hypervisor root passwords must contain only printable ASCII characters.

12. Set up the primary management interface that is used to connect to XenCenter.

If your computer has multiple NICs, select the NIC which you want to use for management. Choose **OK** to proceed.

13. Configure the Management NIC IP address by choosing **Automatic configuration (DHCP)** to configure the NIC using DHCP, or **Static configuration** to configure the NIC manually. To have the management interface on a VLAN network, provide the VLAN ID.

Note:

To be part of a pool, Citrix Hypervisor servers must have static IP addresses or be DNS addressable. When using DHCP, ensure that a static DHCP reservation policy is in place.

14. Specify the hostname and the DNS configuration, manually or automatically via DHCP.

In the **Hostname Configuration** section, select **Automatically set via DHCP** to have the DHCP server provide the hostname along with the IP address. If you select **Manually specify**, enter the hostname for the server in the field provided.

Note:

If you manually specify the hostname, enter a short hostname and *not the fully qualified domain name (FQDN)*. Entering an FQDN can cause external authentication to fail, or the Citrix Hypervisor server might be added to AD with a different name.

In the **DNS Configuration** section, choose **Automatically set via DHCP** to get name service configuration using DHCP. If you select **Manually specify**, enter the IP addresses of your primary (required), secondary (optional), and tertiary (optional) DNS servers in the fields provided.

Select **OK** to proceed.

15. 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.
16. Specify how you want the server to determine local time: using NTP or manual time entry. Make your selection, and choose **OK** to proceed.
 - If using NTP, select **NTP is configured by my DHCP server** or enter at least one NTP server name or IP address in the fields below. Choose **OK** to proceed.
 - If you elected to set the date and time manually, you are prompted to do so. Choose **OK** to proceed.

Note:

Citrix Hypervisor assumes that the time setting in the BIOS of the server is the current time in UTC.

17. Select **Install Citrix Hypervisor**.

The installation process starts. This might take some minutes.

18. The next screen asks if you want to install any supplemental packs (including driver disks). If you plan to install any supplemental packs or driver disks provided by your hardware supplier, choose **Yes**.

Note:

If you have already loaded a driver disk during initial installation, you might be prompted to reinsert the driver disk so that the driver can be installed onto disk. At this point, reinsert the driver disk to ensure that your Citrix Hypervisor instance contains the new driver.

If you choose to install supplemental packs, you are prompted to insert them. Eject the Citrix Hypervisor installation media, and insert the supplemental pack media. Choose **OK**.

Select **Use media** to proceed with the installation.

Repeat for each pack to be installed.

19. From the **Installation Complete** screen, eject the installation media (if installing from USB or CD) and select **OK** to reboot the server.

After the server reboots, Citrix Hypervisor 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 Citrix Hypervisor server.

Install XenCenter

XenCenter must be installed on a Windows machine that can connect to the Citrix Hypervisor server through your network. Ensure that .NET framework version 4.6 or above is installed on this system.

To install XenCenter:

1. Download the installer for the latest version of XenCenter from the [Citrix Hypervisor Download 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 [XenCenter documentation](#).

Connect XenCenter to the Citrix Hypervisor server

To connect XenCenter to the Citrix Hypervisor server:

1. Launch XenCenter. The program opens to the **Home** tab.
2. Click the **Add New Server** icon.
3. Enter the IP address of the Citrix Hypervisor server in the **Server** field. Type the root user name and password that you set during Citrix Hypervisor 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 using XenCenter or the Windows Registry Editor.

To do so in XenCenter: from the 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.

To do so using the Windows Registry Editor, navigate to the key `HKEY_LOCAL_MACHINE\Software\Citrix\XenCenter` and add a key named `AllowCredentialSave` with the string value `true` or `false`.

Installation and deployment scenarios

January 9, 2023

This section steps through the following common installation and deployment scenarios:

- One or more Citrix Hypervisor servers with local storage
- Pools of Citrix Hypervisor servers with shared storage:
 - Multiple Citrix Hypervisor servers with shared NFS storage
 - Multiple Citrix Hypervisor servers with shared iSCSI storage

Citrix Hypervisor servers with local storage

The simplest deployment of Citrix Hypervisor is to run VMs on one or more Citrix Hypervisor servers with local storage.

Note:

Live migration of VMs between Citrix Hypervisor servers is only available when they share storage. However, storage live migration is still available.

Basic hardware requirements

- One or more 64-bit x86 servers with local storage
- One or more Windows systems, on the same network as the Citrix Hypervisor servers

High-level procedure

1. Install the Citrix Hypervisor server software on the servers.
2. Install XenCenter on the Windows systems.
3. Connect XenCenter to the Citrix Hypervisor servers.

After you connect XenCenter to the Citrix Hypervisor servers, storage is automatically configured on the local disk of the hosts.

Pools of Citrix Hypervisor servers with shared storage

A *pool* comprises multiple Citrix Hypervisor server installations, bound together as a single managed entity. When combined with shared storage, a pool enables VMs to be started on *any* Citrix Hypervisor server in the pool that has sufficient memory. The VMs can then dynamically be moved between hosts while running (live migration) with minimal downtime. If an individual Citrix Hypervisor server suffers a hardware failure, you can restart the failed VMs on another host in the same pool.

If the High Availability (HA) feature is enabled, protected VMs are automatically moved if there is a host failure.

To set up *shared storage* between hosts in a pool, create a storage repository. Citrix Hypervisor storage repositories (SR) are storage containers in which virtual disks are stored. SRs, like virtual disks, are persistent, on-disk objects that exist independently of Citrix Hypervisor. SRs can exist on different types of physical storage devices, both internal and external, including local disk devices and shared network storage. Several different types of storage are available when you create an SR, including:

- NFS VHD storage
- Software iSCSI storage
- Hardware HBA storage
- GFS2 storage

This following sections step through setting up two common shared storage solutions—NFS and iSCSI—for a pool of Citrix Hypervisor servers. Before you create an SR, configure your NFS or iSCSI storage. Setup differs depending on the type of storage solution that you use. For details, see your vendor documentation. In all cases, to be part of a pool, the servers providing shared storage must have static IP addresses or be DNS addressable. For further information on setting up shared storage, see [Storage](#).

We recommend that you create a pool before you add shared storage. For pool requirements and setup procedures, see [Pool Requirements](#) in the XenCenter documentation or [Hosts and Resource Pools](#).

Citrix Hypervisor servers with shared NFS storage

Basic hardware requirements

- Two or more 64-bit x86 servers with local storage
- One or more Windows systems, on the same network as the Citrix Hypervisor servers
- A server exporting a shared directory over NFS

High-level procedure

1. Install the Citrix Hypervisor server software on the servers.
2. Install XenCenter on the Windows systems.
3. Connect XenCenter to the Citrix Hypervisor servers.
4. Create your pool of Citrix Hypervisor servers.
5. Configure the NFS server.
6. Create an SR on the NFS share at the pool level.

Configuring your NFS storage

Before you create an SR, configure the NFS storage. To be part of a pool, the NFS share must have a static IP address or be DNS addressable. Configure the NFS server to have one or more targets that can be mounted by NFS clients (for example, Citrix Hypervisor servers in a pool). Setup differs depending on your storage solution, so it is best to see your vendor documentation for details.

To create an SR on the NFS share at the pool level in XenCenter:

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 VHD as the storage type. Choose **Next** to continue.
3. 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 server must be configured to export the specified path to all Citrix Hypervisor servers in the pool.

4. Click **Finish**.

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

Creating an SR on the NFS share at the pool level using the xe CLI

1. Open a console on any Citrix Hypervisor server in the pool.
2. Create the storage repository on *server:/path* by entering the following:

```
1 xe sr-create content-type=user type=nfs name-label=sr_name= \  
2   shared=true device-config:server=server \  
3   device-config:serverpath=path
```

The `device-config-server` argument refers to the name of the NFS server and the `device-config-serverpath` argument refers to the path on the server. Since `shared` is set to true, the shared storage is automatically connected to every host in the pool. Any hosts that later join are also connected to the storage. The UUID of the created storage repository is printed to the console.

3. Find the UUID of the pool by using the `pool-list` command.
4. Set the new SR as the pool-wide default by entering the following:

```
1 xe pool-param-set uuid=pool_uuid \  
2   default-SR=storage_repository_uuid
```

As shared storage has been set as the pool-wide default, all future VMs have their disks created on this SR.

Citrix Hypervisor servers with shared iSCSI storage**Basic hardware requirements**

- Two or more 64-bit x86 servers with local storage
- One or more Windows systems, on the same network as the Citrix Hypervisor servers
- A server providing a shared directory over iSCSI

High-level procedure

1. Install the Citrix Hypervisor server software on the servers.
2. Install XenCenter on the Windows systems.
3. Connect XenCenter to the Citrix Hypervisor servers.
4. Create your pool of Citrix Hypervisor servers.
5. Configure the iSCSI storage.
6. If necessary, enable multiple initiators on your iSCSI device.

7. If necessary, configure the iSCSI Qualified Name (IQN) for each Citrix Hypervisor server.
8. Create an SR on the iSCSI share at the pool level.

Configuring your iSCSI storage

Before you create an SR, configure the iSCSI storage. To be part of a pool, the iSCSI storage must have a static IP address or be DNS addressable. Provide an iSCSI target LUN on the SAN for the VM storage. Configure Citrix Hypervisor servers to be able to see and access the iSCSI target LUN. Both the iSCSI target and each iSCSI initiator on each Citrix Hypervisor server must have a valid and **unique** IQN. For configuration details, it is best to see your vendor documentation.

Configuring an iSCSI IQN for each Citrix Hypervisor server

Upon installation, Citrix Hypervisor automatically attributes a unique IQN to each host. If you must adhere to a local administrative naming policy, you can change the IQN by using the following xe CLI command:

```
1 xe host-param-set uuid=<host_uuid> iscsi_iqn=<iscsi_iqn>
```

To create an SR on the iSCSI share at the pool level using XenCenter:

Warning:

When you create Citrix Hypervisor SRs on iSCSI or HBA storage, any existing contents of the volume are destroyed.

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 Software iSCSI as the storage type. Choose **Next** to continue.
3. 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 Citrix Hypervisor server in the pool to have access to one or more LUNs.

4. If you have configured the iSCSI target to use CHAP authentication, enter the User and Password.
5. Click the **Discover IQNs** button, and then choose the iSCSI target IQN from the Target IQN list.

Warning:

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

6. Click the **Discover LUNs** button, and then choose 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.

7. Click **Finish**.

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

To create an SR on the iSCSI share at the pool level by using the xe CLI:**Warning:**

When you create Citrix Hypervisor SRs on iSCSI or HBA storage, any existing contents of the volume are destroyed.

1. On the console of any server in the pool, run the command:

```
1 xe sr-create name=label=name_for_sr \  
2   host-uuid=host_uuid device-config:target=  
3   iscsi_server_ip_address \  
4   device-config:targetIQN=iscsi_target_iqn device-config:SCSIid=  
   scsi_id \  
   content-type=user type=lvmoiscsi shared=true
```

The `device-config:target` argument refers to the name or IP address of the iSCSI server. Since the `shared` argument is set to **true**, the shared storage is automatically connected to every host in the pool. Any hosts that later join are also connected to the storage.

The command returns the UUID of the created storage repository.

2. Find the UUID of the pool by running the `pool-list` command.
3. Set the new SR as the pool-wide default by entering the following:

```
1 xe pool-param-set uuid=pool_uuid default-SR=iscsi_shared_sr_uuid
```

As shared storage has been set as the pool-wide default, all future VMs have their disks created on this SR.

Upgrade from an existing version

June 11, 2024

This article describes how to upgrade Citrix Hypervisor by using XenCenter or the xe CLI. It guides you through upgrading your Citrix Hypervisor servers—both pooled and standalone—automatically (using the XenCenter Rolling Pool Upgrade wizard) and manually.

Important:

There are no supported upgrade paths to Citrix Hypervisor 8.2 Cumulative Update 1. All earlier versions of XenServer and Citrix Hypervisor are out of support. If you attempt to upgrade to Citrix Hypervisor 8.2 Cumulative Update 1 from an out-of-support version you are not supported and do so at your own risk.

This article is provided for reference only.

For out-of-support versions of XenServer and Citrix Hypervisor you cannot upgrade to Citrix Hypervisor 8.2 Cumulative Update 1 directly. Perform a clean installation using the **Base Installation ISO**. For more information, see [Install](#).

To retain VMs from your previous installation of Citrix Hypervisor or XenServer, when no upgrade path is available, export the VMs and import them into your clean installation of Citrix Hypervisor 8.2 Cumulative Update 1. VMs exported from any supported version of Citrix Hypervisor or XenServer can be imported into Citrix Hypervisor 8.2 Cumulative Update 1. For more information, see [Import and export VMs](#).

If you are a Citrix Hypervisor 8.2 Cumulative Update 1 customer who wants to upgrade to a newer version of XenServer, see [Upgrade from Citrix Hypervisor 8.2 Cumulative Update 1](#) in the XenServer 8 documentation.

Before you start

Review the following information before starting your upgrade. Take the necessary steps to ensure that your upgrade process is successful.

- Upgrading Citrix Hypervisor servers, and particularly a pool of Citrix Hypervisor servers, requires careful planning and attention. To avoid losing any existing data, either:
 - Map your upgrade path carefully.
 - Use the XenCenter Rolling Pool Upgrade wizard, and ensure that you select the option to *upgrade* when you are stepping through the installer.

- Check that the hardware your pool is installed on is compatible with the version of Citrix Hypervisor you are about to upgrade to. For more information, see the [Hardware Compatibility List \(HCL\)](#).
- If you are using XenCenter to upgrade your hosts, download and install the latest version of XenCenter from the [Citrix Hypervisor download site](#).

For example, when upgrading to Citrix Hypervisor 8.2, use the latest version of XenCenter issued for Citrix Hypervisor 8.2. Using earlier versions of XenCenter to upgrade to a newer version of Citrix Hypervisor is not supported.

- Check that the operating systems of your VMs are supported by the version of Citrix Hypervisor you are about to upgrade to. If your VM operating system is not supported in the target version of Citrix Hypervisor, upgrade your VM operating system to a supported version. For more information, see [Guest operating system support](#).
- Paravirtualized (PV) VMs are not supported in Citrix Hypervisor 8.2 Cumulative Update 1. 32-bit PV VMs are blocked from starting on Citrix Hypervisor 8.2 Cumulative Update 1 servers. Ensure that before upgrading you remove any PV VMs from your pool or upgrade your VMs to a supported version of their operating system. For more information, see [Upgrade from PV to HVM guests](#).

Earlier versions of the Citrix License Server virtual appliance run in PV mode. We recommend that you transition to using the Windows-based Citrix License Server as part of upgrading to Citrix Hypervisor 8.2 Cumulative Update 1.

- If you have Windows VMs running in your pool that will be migrated as part of your upgrade, 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
- If you have Linux VMs running in your pool that will be migrated as part of your upgrade, ensure that the latest version of the Citrix VM Tools for Linux is installed.
- Boot-from-SAN settings are *not* inherited during the manual upgrade process. When upgrading using the ISO or PXE process, follow the same instructions as used in the installation process below to ensure that `multipathd` is correctly configured. For more information, see [Boot from SAN](#).
- Quiesced snapshots are no longer supported. If you have existing snapshot schedules that create quiesced snapshots, these snapshot schedules fail after the upgrade. To ensure that snapshots continue to be created, delete the existing schedule and create a new one that creates non-quiesced snapshots before performing the upgrade.
- Legacy SSL mode is no longer supported. Disable this mode on all hosts in your pool before attempting to upgrade to the latest version on Citrix Hypervisor. To disable legacy SSL mode,

run the following command on your pool master before you begin the upgrade: `xe pool-disable-ssl-legacy uuid=<pool_uuid>`

- The Container Management supplemental pack is no longer supported. After you update or upgrade to the latest version of Citrix Hypervisor, you can no longer use the features of this supplemental pack.
- When you upgrade Citrix Hypervisor, previously applied supplemental packs are removed and so they must be reapplied during or after the upgrade.
- The vSwitch Controller is no longer supported. Disconnect the vSwitch Controller from your pool before attempting to upgrade to the latest version on Citrix Hypervisor. After the upgrade, the following configuration changes take place:
 - Cross-server private networks revert to single-server private networks.
 - Any Quality of Service settings made through the DVSC console are no longer applied. Network rate limits are no longer enforced.
 - ACL rules are removed. All traffic from VMs is allowed.
 - Port mirroring (RSPAN) is disabled.

If, after update or upgrade, you find leftover state about the vSwitch Controller in your pool, clear the state with the following CLI command: `xe pool-set-vswitch-controller address=`

Rolling pool upgrades

Citrix Hypervisor enables you to perform a rolling pool upgrade. A rolling pool upgrade keeps all the services and resources offered by the pool available while upgrading all of the hosts in a pool. This upgrade method takes only one Citrix Hypervisor server offline at a time. Critical VMs are kept running during the upgrade process by live migrating the VMs to other hosts in the pool.

Note:

The pool must have shared storage to keep your VMs running during a rolling pool upgrade. If your pool does not have shared storage, you must stop your VMs before upgrading because the VMs cannot be live migrated.

Storage live migration is not supported with rolling pool upgrades.

You can perform a rolling pool upgrade using XenCenter or the xe CLI. When using XenCenter, we recommend using the Rolling Pool Upgrade wizard. This wizard organizes the upgrade path automatically and guides you through the upgrade procedure. If you are using the xe CLI, first plan your upgrade path and then live migrate running VMs between Citrix Hypervisor servers as you perform the rolling pool upgrade manually.

The Rolling Pool Upgrade wizard is available for licensed Citrix Hypervisor customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, see [Licensing](#). To upgrade, or to get a Citrix Hypervisor license, visit the [Citrix website](#).

Important:

Do not use Rolling Pool Upgrade with Boot from SAN environments. For more information on upgrading boot from SAN environments, see [Boot from SAN](#).

Upgrade Citrix Hypervisor servers by using the XenCenter Rolling Pool Upgrade wizard

The Rolling Pool Upgrade wizard enables you to upgrade Citrix Hypervisor servers, hosts in a pool or standalone hosts, to the current version of Citrix Hypervisor.

The Rolling Pool Upgrade wizard guides you through the upgrade procedure and organizes the upgrade path automatically. For pools, each of the hosts in the pool is upgraded in turn, starting with the pool master. Before starting an upgrade, the wizard conducts a series of prechecks. These prechecks ensure certain pool-wide features, such as high availability are temporarily disabled and that each host in the pool is prepared for upgrade. Only one host is offline at a time. Any running VMs are automatically migrated off each host before the upgrade is installed on that host.

The Rolling Pool Upgrade wizard also allows you to automatically apply the available hotfixes when upgrading to a newer version of Citrix Hypervisor. This enables you to bring your standalone hosts or pools up-to-date with a minimum number of reboots at the end. You must be connected to the Internet during the upgrade process for this feature to work.

You can benefit from the automatic application of hotfixes feature when you use XenCenter issued with Citrix Hypervisor 8.2 Cumulative Update 1 to upgrade from any supported version of Citrix Hypervisor or XenServer.

Note:

Rolling Pool Upgrade using XenCenter is only available for licensed Citrix Hypervisor customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.

The wizard can operate in **Manual** or **Automatic** mode:

- In **Manual Mode**, you must manually run the Citrix Hypervisor installer on each host in turn and follow the on-screen instructions on the serial console of the host. When the upgrade begins, XenCenter prompts you to insert the XenCenter installation media or specify a network boot server for each host that you upgrade.

- In **Automatic Mode**, the wizard uses network installation files on an HTTP, NFS, or FTP server to upgrade each host in turn. This mode doesn't require you to insert installation media, manually reboot, or step through the installer on each host. 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.

Note:

If you are using IIS to host the installation media, ensure that double escaping is enabled on IIS before extracting the installation ISO on it.

Before you upgrade

Before you begin your upgrade, be sure to make the following preparations:

- Download and install the latest version of XenCenter provided for Citrix Hypervisor 8.2 Cumulative Update 1 from the [Citrix Hypervisor Product Download](#) page. Using earlier versions of XenCenter to upgrade to a newer version of Citrix Hypervisor is not supported.
- We strongly recommend that you take a backup of the state of your existing pool using the `pool -dump-database` xe CLI command. For more information, see [Command line interface](#). Taking a backup state ensures that you can revert a partially complete rolling upgrade to its original state without losing VM data.
- Ensure that your hosts are not over-provisioned: check that hosts have sufficient memory to carry out the upgrade.

As a general guideline, if N equals the total number of hosts in a pool, there must be sufficient memory across N-1 hosts to run all of the live VMs in the pool. It is best to suspend any non-critical VMs during the upgrade process.

- If you have vGPU-enabled VMs running on your pool, complete the following steps to migrate the pool while these VMs are running:
 - Ensure that the GPU you are using is supported on the version you plan to upgrade to.
 - Identify a version of the NVIDIA drivers that is available for both your current version of Citrix Hypervisor and the version of Citrix Hypervisor you are upgrading. If possible, choose the latest available drivers.
 - Install the new NVIDIA drivers on your Citrix Hypervisor servers and the matching guest drivers on any of your vGPU-enabled VMs.
 - Ensure that you also have the version of the NVIDIA driver that matches the version of Citrix Hypervisor that you are upgrading to. You are prompted to install these drivers as a supplemental pack as part of the Rolling Pool Upgrade process.

Rolling Pool Upgrade wizard checks that the following actions have been taken. Perform these actions before you begin the upgrade process:

- Empty the CD/DVD drives of the VMs in the pools.
- Disable high availability.

Upgrade process

To upgrade Citrix Hypervisor hosts by using the XenCenter Rolling Pool Upgrade wizard:

1. Open the Rolling Pool Upgrade wizard: on the **Tools** menu, select **Rolling Pool Upgrade**.
2. Read the **Before You Start** information, and then 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
 - **Manual Mode** for a manual upgrade from either a USB/CD/DVD or using network boot (using existing infrastructure)

Notes:

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.

If you choose **Manual Mode**, you must run the Citrix Hypervisor 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 Citrix Hypervisor installation media or specify a network boot server for each host that you upgrade.

5. Choose whether you want XenCenter to automatically download and install the minimal set of updates (hotfixes) after upgrading the servers to a newer version. The apply updates option is selected by default. However, you must have an internet connection to download and install the updates.
6. After you have selected your Upgrade Mode, click **Run Prechecks**.
7. Follow the recommendations to resolve any upgrade prechecks that have failed. If you want XenCenter to resolve all failed prechecks automatically, click **Resolve All**.

When all prechecks have been resolved, click **Next** to continue.

8. Prepare the Citrix Hypervisor 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 Citrix Hypervisor pool.
- Citrix Hypervisor supports FTP in passive mode only.

If you chose **Manual Mode**, note the upgrade plan and instructions.

Click **Start Upgrade**.

9. 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.

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 Citrix Hypervisor version you are upgrading to.

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.

10. The Rolling Pool Upgrade wizard prints a summary when the upgrade is complete. Click **Finish** to close the wizard.

Notes:

After a Rolling Pool Upgrade is complete, a VM might not be located on its home server. To relocate the VM, you can do one of the following actions:

- Live migrate the VM to its home server
- Shut down the VM and then start it on its home server

Upgrade Citrix Hypervisor servers by using the xe CLI

Performing a rolling pool upgrade using the xe CLI requires careful planning. Be sure to read the following section with care before you begin.

Plan an upgrade path

As you plan your upgrade, it is important to be aware of the following:

- You can only migrate VMs from Citrix Hypervisor servers running an older version of Citrix Hypervisor to one running the same version or higher. For example, from version 7.0 to version 7.1 Cumulative Update 2 or from version 7.1 Cumulative Update 2 to version 8.2 Cumulative Update 1.

You **cannot** migrate VMs from an upgraded host to one running an older version of Citrix Hypervisor. For example, from version 8.2 Cumulative Update 1 to version 7.1 Cumulative Update 2. Be sure to allow for space on your Citrix Hypervisor servers accordingly.

- We strongly advise against running a mixed-mode pool (one with multiple versions of Citrix Hypervisor co-existing) for longer than necessary, as the pool operates in a degraded state during upgrade.
- Key control operations are not available during the upgrade. Do not attempt to perform any control operations. Though VMs continue to function as normal, VM actions other than migrate 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.
- Always upgrade the master host first. Do not place the host into maintenance mode using XenCenter before performing the upgrade. If you put the master in maintenance mode, a new master is designated.
- After upgrading a host, apply any hotfixes that have been released for the upgraded version of Citrix Hypervisor before migrating VMs onto the host.
- We strongly recommend that you take a backup of the state of your existing pool using the `pool-dump-database` CLI command. For more information, see [Command Line interface](#). This allows you to revert a partially complete rolling upgrade back to its original state without losing any VM data. If you have to revert the rolling upgrade for any reason, you might have to shut down VMs. This action is required because it is not possible to migrate a VM from an upgraded Citrix Hypervisor server to a host running an older version of Citrix Hypervisor.

Before you begin your rolling pool upgrade

- If you are using XenCenter, upgrade XenCenter to the latest version provided on the [Citrix download site](#). The newer version of XenCenter correctly controls older versions of Citrix Hypervisor servers.
- Empty the CD/DVD drives of the VMs in the pool. For details and instructions, see *Before Upgrading a Single Citrix Hypervisor server*.
- Disable high availability.

Perform rolling pool upgrades by using the xe CLI

1. **Start with the pool master.** Disable the master by using the `host-disable` command. This prevents any new VMs from starting on the specified host.
2. Ensure that no VMs are running on the master. Shut down, suspend or migrate VMs to other hosts in the pool.

To migrate specified VMs to specified hosts, use the `vm-migrate` command. By using the `vm-migrate` command, you have full control over the distribution of migrated VMs to other hosts in the pool.

To live migrate all VMs to other hosts in the pool, use the `host-evacuate` command. By using the `host-evacuate` command, you leave the distribution of migrated VMs to Citrix Hypervisor.

3. Shut down the pool master.

Important:

You are unable to contact the pool master until the upgrade of the master is complete. Shutting down the pool master causes the other hosts in the pool to enter *emergency mode*. Hosts can enter emergency mode when they are in a pool whose master 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 master using the Citrix Hypervisor installation media and method of your choice (such as, USB or network). Follow the Citrix Hypervisor installation procedure until the installer offers you the option to upgrade. **Choose to upgrade.** For more information, see [Install](#).

Warnings:

- Ensure you select the upgrade option to avoid losing any existing data.
- If anything interrupts the upgrade of the pool master or if the upgrade fails for any reason, do not attempt to proceed with the upgrade. Reboot the pool master and restore to a working version of the master.

When your pool master restarts, the other hosts in the pool leave emergency mode and normal service is restored after a few minutes.

5. Apply any hotfixes that have been released for the new version of Citrix Hypervisor to the pool master.
6. On the pool master, start or resume any shutdown or suspended VMs. Migrate any VMs that you want back to the pool master.
7. Select the next Citrix Hypervisor server in your upgrade path. Disable the host.

8. Ensure that no VMs are running on the host. Shut down, suspend or migrate VMs to other hosts in the pool.
9. Shut down the host.
10. Follow the upgrade procedure for the host, as described for the master in Step 4.

Note:

If the upgrade of a host that is not the master fails or is interrupted, you do not have to revert. Use the `host-forget` command to forget the host. Reinstall Citrix Hypervisor on the host, and then join it, as a new host, to the pool using the `pool-join` command.

11. Apply any hotfixes that have been released for the new version of Citrix Hypervisor to the host.
12. On the host, start or resume any shutdown or suspended VMs. Migrate any VMs that you want back to the host.
13. Repeat Steps 6–10 for the rest of the hosts in the pool.

Upgrade a single Citrix Hypervisor server by using the xe CLI

Before you upgrade a single Citrix Hypervisor server

Before upgrading a standalone Citrix Hypervisor server, shut down or suspend any VMs running on that host. It is important to eject and empty CD/DVD drives of any VMs you plan to suspend. If you do not empty the CD/DVD drives, you may not be able to resume the suspended VMs after upgrade.

An *empty* VM CD/DVD drive means the VM is not attached to an ISO image or a physical CD/DVD mounted through the Citrix Hypervisor server. In addition, you must ensure that the VM is not attached to any physical CD/DVD drive on the Citrix Hypervisor server at all.

To empty the CD/DVD drive of a VM by using the xe CLI:

1. Identify which VMs do not have empty CD/DVD drives by typing the following:

```
1 xe vbd-list type=CD empty=false
```

This returns a list of all the VM CD/DVD drives that are not empty, for example:

```
1      uuid ( RO) : abae3997-39af-2764-04a1-ffc501d132d9
2      vm-uuid ( RO): 340a8b49-866e-b27c-99d1-fb41457344d9
3      vm-name-label ( RO): VM02_DemoLinux
4      vdi-uuid ( RO): a14b0345-b20a-4027-a233-7cbd1e005ede
5      empty ( RO): false
6      device ( RO): xvdd
7
8      uuid ( RO) : ec174a21-452f-7fd8-c02b-86370fa0f654
9      vm-uuid ( RO): db80f319-016d-0e5f-d8db-3a6565256c71
```

```
10    vm-name-label ( R0): VM01_DemoLinux
11    vdi-uuid ( R0): a14b0345-b20a-4027-a233-7cbd1e005ede
12    empty ( R0): false
13    device ( R0): xvdd
```

Note the `uuid`, which is the first item in the list.

2. To empty the CD/DVD drives of the VMs listed, type the following:

```
1  xe vbd-eject uuid=uuid
```

Upgrade a single Citrix Hypervisor server by using the xe CLI

To upgrade a single Citrix Hypervisor server by using the xe CLI:

1. Disable the Citrix Hypervisor server that you want to upgrade by typing the following:

```
1  xe host-disable host-selector=host_selector_value
```

When the Citrix Hypervisor server is disabled, VMs cannot be created or started on that host. VMs also cannot be migrated to a disabled host.

2. Shut down or suspend any VMs running on the host that you want to upgrade by using the `xe vm-shutdown` or `xe vm-suspend` commands.
3. Shut down the host by using the `xe host-shutdown` command.
4. Follow the Citrix Hypervisor installation procedure until the installer offers you the option to upgrade. **Choose to upgrade.** For more information, see [Install](#).

Warning:

Be sure to select the upgrade option to avoid losing any existing data.

You don't have to configure any settings again during the setup procedure. The upgrade process follows the first-time installation process but several setup steps are bypassed. The existing settings for networking configuration, system time, and so on, are retained.

When your host restarts, normal service is restored after a few minutes.

5. Apply any hotfixes that have been released for the new version of Citrix Hypervisor.
6. Restart any shutdown VMs, and resume any suspended VMs.

Update your hosts

July 22, 2024

Updates can often be applied with minimal service interruption. We recommend that customers use XenCenter to apply all updates. If you are updating a Citrix Hypervisor pool, you can avoid VM downtime by using the **Install Update** wizard in XenCenter. The **Install Update** wizard applies updates, updating one host at a time, automatically migrating VMs away from each host as the hotfix or update is applied.

You can configure XenCenter to check periodically for available Citrix Hypervisor and XenCenter updates and new versions. Any Alerts are displayed in the **Notifications** pane.

Note:

Ensure that you use the latest version of XenCenter to apply updates to your Citrix Hypervisor hosts and pools. The latest version of XenCenter is provided on the [Citrix download site](#).

Types of update

The following types of updates are available for Citrix Hypervisor:

- **Releases**, which are full releases of Citrix Hypervisor that can be applied as updates to the supported versions of Citrix Hypervisor.
- **Hotfixes**, which generally supply bug fixes to one or more specific issues. Hotfixes are provided for supported Citrix Hypervisor releases.
- **Cumulative Updates**, which contain previously released hotfixes and can contain support for new guests and hardware. Cumulative updates are applied to Citrix Hypervisor releases from the Long Term Service Release (LTSR) stream.
- **Supplemental packs**, which are provided by our partners and can also be applied as updates to Citrix Hypervisor.
- **Driver disks**, which are a type of supplemental pack that enables you to use the latest hardware.

Notes:

- If you use XenCenter to update your hosts, you must update your XenCenter installation to the latest version before beginning.
- Always update the pool master before updating any other hosts in a pool.

Releases

Citrix Hypervisor 8.2 Cumulative Update 1 is an update for Citrix Hypervisor 8.2. However, now that Citrix Hypervisor 8.2 is no longer supported, updating from Citrix Hypervisor 8.2 to Hypervisor 8.2 Cumulative Update 1 is no longer supported.

Hotfixes

We might release hotfixes for Citrix Hypervisor 8.2 Cumulative Update 1 that provide fixes for specific issues.

Hotfixes for Citrix Hypervisor 8.2 Cumulative Update 1 are made available from the [Citrix Knowledge Center](#). We recommend that customers regularly check the Knowledge Center for new updates. Alternatively, you can subscribe to email alerts for updates to Citrix Hypervisor by registering for an account at <http://www.citrix.com/support/>.

Hotfixes on the latest release are available to all Citrix Hypervisor customers. However, hotfixes on previous releases that are still in support are only available for customers with an active Citrix Customer Success Services (CSS) account.

Hotfixes on the LTSR stream are available to customers with an active CSS account. For more information, See [Licensing](#).

Cumulative Updates

Cumulative Updates are provided for LTSRs of Citrix Hypervisor. These updates provide fixes for issues, and may contain support for new guests and hardware.

Cumulative Updates are available to customers with an active CSS account.

Driver disks

You can install a driver disk using one of the following methods:

- By using XenCenter (recommended)
- During a clean Citrix Hypervisor installation
- By using the xe CLI

For information on how to install a driver disk by using XenCenter, see [Install driver disks](#). For information on how to install a driver disk during a clean Citrix Hypervisor installation, see [Install the Citrix Hypervisor server](#).

After installing the driver, restart your server for the new version of the driver to take effect. As with any software update, we advise you to back up your data before installing a driver disk.

Install a driver disk by using the xe CLI Perform the following steps to install the driver disk remotely using the xe CLI:

1. Download the driver disk to a known location on a computer that has the remote xe CLI installed.

2. Extract the contents of the zip file.

For the next step, ensure that you use the driver ISO and not the ISO that contains the source files.

3. Upload the driver disk:

```
1 xe [connection_parameters] update-upload file-name=
```

The UUID of the driver disk is returned when the upload completes.

4. Apply the driver disk:

```
1 xe [connection_parameters] update-apply uuid=
```

5. To complete the installation, restart the host. The driver does not take effect until after the host is restarted.

Configure XenCenter to download updates

Downloading updates from <https://support.citrix.com> is restricted to customers with a Citrix account. Some updates are only available to customers who are part of Citrix Success Services. These restrictions are now enforced by XenCenter.

To receive updates through XenCenter, you must first install the latest version of XenCenter and obtain a client ID JSON file. For more information, see [Authenticating your XenCenter to receive updates](#).

Prepare a pool for an update

Updates to Citrix Hypervisor can be delivered as a hotfix or a Cumulative Update or a Current Release. Pay careful attention to the release notes published with each update. Each update can have unique installation instructions, particularly regarding preparatory and post-update operations. The following sections offer general guidance and instructions for applying updates to your Citrix Hypervisor systems.

Before you apply an update to the Citrix Hypervisor pool, pay careful attention to the following:

- All hosts in the pool must be running Citrix Hypervisor 8.2 before you apply the hotfix.
- Back up your data before applying an update. For backup procedures, see [Disaster recovery and backup](#).
- Before applying a Cumulative Update, check that the hardware your pool is installed on is compatible with the version of Citrix Hypervisor you are about to update to. For more information, see the [Hardware Compatibility List \(HCL\)](#).

- Before applying a Cumulative Update, check that the operating systems of your VMs are supported by the version of Citrix Hypervisor you are about to update to. If your VM operating system is not supported in the target version of Citrix Hypervisor, upgrade your VM operating system to a supported version. For more information, see [Guest operating system support](#).
- Paravirtualized (PV) VMs are not supported in Citrix Hypervisor 8.2 Cumulative Update 1. 32-bit PV VMs are blocked from starting on Citrix Hypervisor 8.2 Cumulative Update 1 servers. Ensure that before updating you remove any PV VMs from your pool or upgrade your VMs to a supported version of their operating system. For more information, see [Upgrade from PV to HVM guests](#).

Earlier versions of the Citrix License Server virtual appliance run in PV mode. We recommend that you transition to using the Windows-based Citrix License Server as part of updating to Citrix Hypervisor 8.2 Cumulative Update 1.

- If you have Windows VMs running in your pool that will be migrated as part of the update, take the following steps for each VM:
 - Set the value of the following registry key to a REG_DWORD value of '3': `HLKM\System\CurrentControlSet\services\xenbus_monitor\Parameters\Autoreboot`
 - Ensure that the latest version of the XenServer VM Tools for Windows is installed
 - Take a snapshot of the VM
- If you have Linux VMs running in your pool that will be migrated as part of your upgrade, ensure that the latest version of the Citrix VM Tools for Linux is installed.
- Update all servers in a pool within a short period: running a mixed-mode pool (a pool that includes updated and non-updated servers) is not a supported configuration. Schedule your updates to minimize the amount of time that a pool runs in a mixed state.
- Update all servers within a pool sequentially, always starting with the pool master. XenCenter's **Install Update** wizard manages this process automatically.
- After applying an update to all hosts in a pool, update any required driver disks before restarting Citrix Hypervisor servers.
- After applying a Cumulative Update or Current Release to a host, apply any hotfixes released for that Cumulative Update or Current Release before migrating VMs onto the host.
- Legacy SSL mode is no longer supported. Disable this mode on all hosts in your pool before attempting to update to the latest version on Citrix Hypervisor. To disable legacy SSL mode, run the following command on your pool master before you begin the update: `xe pool-disable-ssl-legacy uuid=<pool_uuid>`
- The Container Management supplemental pack is no longer supported. After you update or upgrade to the latest version of Citrix Hypervisor, you can no longer use the features of this supplemental pack.

- The vSwitch Controller is no longer supported. Disconnect the vSwitch Controller from your pool before attempting to update to the latest version on Citrix Hypervisor. After the update, the following configuration changes take place:
 - Cross-server private networks revert to single-server private networks.
 - Any Quality of Service settings made through the DVSC console are no longer applied. Network rate limits are no longer enforced.
 - ACL rules are removed. All traffic from VMs is allowed.
 - Port mirroring (RSPAN) is disabled.

After update or upgrade, if you find leftover state about the vSwitch Controller in your pool, clear the state with the following CLI command: `xe pool-set-vswitch-controller address=`

Before you begin updating

- Log into a user account with full access permissions (for example, as a Pool Administrator or using a local root account).
- Empty the CD/DVD drives of any VMs you plan to suspend. For details and instructions, see [Before Upgrading a Single Citrix Hypervisor server](#).
- If applicable, disable high availability.

Apply updates to a pool

The update installation mechanism in XenCenter allows you to download and extract the selected update from the Support website. You can apply an update to multiple hosts and pools simultaneously using the **Install Update** wizard. During the process, the **Install Update** wizard completes the following steps for each server:

- Migrates VMs off the server
- Places the server in maintenance mode
- Applies the update to the server
- Reboots the host if necessary
- Migrates the VMs back to the updated host.

Any actions taken at the precheck stage to enable the updates to be applied, such as turning off HA, are reverted.

The **Install Update** wizard carries out a series of checks known as Prechecks before starting the update process. These checks ensure that the pool is in a valid configuration state. It then manages the update path and VM migration automatically.

Apply updates automatically

XenCenter allows you to apply automated updates that are required to bring your servers up-to-date. You can apply these updates to one or more pools. When you apply automated updates, XenCenter applies the minimum set of updates that are required to bring the selected pool or the standalone server up-to-date. XenCenter minimizes the number of reboots required to bring the pool or the standalone server pool up-to-date. Where possible, XenCenter limits it to a single reboot at the end. For more information, see [Apply Automated Updates](#).

View available updates

The **Updates** section of the **Notifications** view lists the updates that are available for all connected servers and pools.

Notes:

- By default, XenCenter periodically checks for Citrix Hypervisor and XenCenter updates. Click **Refresh** to check manually for available updates.
- If you have disabled automatic check for updates, a message appears on the **Updates** tab. Click **Check for Updates Now** to check for updates manually.

You can select from the **View** list whether to view the list of updates **By Update** or **By Server**.

When you view the list of updates by update, XenCenter displays the list of updates. You can order these updates by server/pool or by date.

- Cumulative Updates and new releases are displayed at the top of this list. Not all new releases can be applied as an update.
- To export this information as a .csv file, click **Export All**. The .csv file lists the following information:
 - Update name
 - Description of the update
 - Servers that this update can be applied to
 - Timestamp of the update
 - A reference to the webpage that the update is downloaded from
- To apply an update to a server, from the **Actions** list for that update select **Download and Install**. This option extracts the update and opens the **Install Update** wizard on the **Select Servers** page with the relevant servers selected. For more information, see [Apply an update to a pool](#).
- To open the release note of an update in your browser, click the **Actions** list and select **Go to Web Page**.

When you view the list of updates by server, XenCenter displays the list of servers connected to XenCenter. This list shows both the updates that you can apply to the servers and the updates that are already installed on the servers.

- To export this information as a .csv file, click **Export All**. The .csv file lists the following information:
 - Pool that the server belongs to
 - Server name
 - Status of the installed Citrix Hypervisor
 - Update status of the server
 - Required updates for this server
 - Installed updates for this server.
- To apply the updates, click **Install Updates**. This choice opens the **Install Update** wizard on the **Select Update** page. For more information, see [Apply an update to a pool](#).

Apply an update to a pool

To apply an update to a pool by using XenCenter:

1. From the XenCenter menu, select **Tools** and then **Install Update**.
2. Read the information displayed on the **Before You Start** page and then click **Next**.
3. The Install Update wizard lists available updates on the **Select Update** page. Select the required update from the list and then click **Next**.
4. On the **Select Servers** page, select the pool and servers that you want to update.

When applying a Cumulative Update or a Current Release, you can also select whether to apply the minimal set of hotfixes for the CU or CR.

Click **Next**.

5. The **Install Update** wizard performs several prechecks to ensure that the pool is in a valid configuration state.

The wizard also checks the following conditions:

- Whether the hosts must be rebooted after the update is applied and displays the result.
 - Whether a live patch is available for the hotfix and whether the live patch can be applied to the hosts. For information about live patching, see [Live Patching](#).
6. Follow the on-screen recommendations to resolve any update prechecks that have failed. If you want XenCenter to resolve all failed prechecks automatically, click **Resolve All**. When the prechecks have been resolved, click **Next**.

7. If you are installing a CU or a CR, XenCenter downloads the updates, uploads them to the default SR of the pool, and installs the updates. The **Upload and Install** page displays the progress.

Notes:

- If the default SR in a pool is not shared or does not have enough space, XenCenter tries to upload the update to another shared SR. If none of the shared SRs have sufficient space, the update is uploaded to local storage of the pool master.
- If the update process cannot complete for any reason, XenCenter halts the process. This action allows you to fix the issue and resume the update process by clicking the **Retry** button.

See Step 10. to complete the installation process.

8. If you are installing a hotfix, choose an **Update Mode**. Review the information displayed on the screen and select an appropriate mode. If the hotfix contains a live patch that can be successfully applied to the hosts, it displays **No action required** on the **Tasks to be performed** screen.

Note:

If you click **Cancel** at this stage, the Install Update wizard reverts the changes and removes the update file from the server.

9. Click **Install update** to proceed with the installation. The Install Update wizard shows the progress of the update, displaying the major operations that XenCenter performs while updating each server in the pool.
10. When the update is applied, click **Finish** to close Install Update wizard. If you chose to perform post-update tasks manually, do so now.

Update a pool of Citrix Hypervisor servers by using the xe CLI

Ensure that you update the pool master before you update any other pool member.

To update a pool of Citrix Hypervisor hosts by using the xe CLI:

1. Download the update file to a known location on the computer running the xe CLI. Note the path to the file.
2. Upload the update file to the pool you want to update by running the following:

```
1 xe -s server -u username -pw password update-upload file-name=  
   filename [sr-uuid=storage_repository_uuid]
```

Here, **-s** refers to the name of the pool master. Citrix Hypervisor assigns the update file a UUID, which this command prints. Note the UUID.

Tip:

After an update file has been uploaded to the Citrix Hypervisor server, you can use the `update-list` and `update-param-list` commands to view information about the file.

3. If Citrix Hypervisor detects any errors or preparatory steps that have not been taken, it alerts you. Be sure to follow any guidance before continuing with the update.

If necessary, you can shut down or suspend any VMs on the hosts that you want to update by using the `vm-shutdown` or `vm-suspend` commands.

To migrate specified VMs to specified hosts, use the `vm-migrate` command. By using the `vm-migrate` command, you have full control over the distribution of migrated VMs to other hosts in the pool.

To live migrate all VMs to other hosts in the pool automatically, use the `host-evacuate` command. By using the `host-evacuate` command, you leave the distribution of migrated VMs to Citrix Hypervisor.

4. Update the pool, specifying the UUID of the update file, by running the following:

```
1 xe update-pool-apply uuid=UUID_of_file
```

This command applies the update or hotfix to all hosts in the pool, starting with the pool master.

Or, to update and restart hosts in a rolling manner, you can apply the update file to an individual host by running the following command:

```
1 xe update-apply host=host uuid=UUID_of_file
```

5. Verify that the update was applied by using the `update-list` command. If the update has been successful, the `hosts` field contains the host UUID.
6. Perform any post-update operations that are required, such as restarting the XAPI toolstack or rebooting the hosts. Perform these operations on the pool master first.

Ensure that you apply the update to all hosts in the pool. Running a mixed-mode pool (a pool that includes updated and non-updated servers) is not a supported configuration.

Update individual hosts by using the xe CLI

To update individual hosts by using the xe CLI:

1. Download the update file to a known location on the computer running the xe CLI. Note the path to the file.

2. Shut down or suspend any VMs on the hosts that you want to update by using the `vm-shutdown` or `vm-suspend` commands.
3. Upload the update file to the host you want to update by running the following:

```
1 xe -s server -u username -pw password update-upload file-name=
  filename [sr-uuid=storage_repository_uuid]
```

Here, `-s` refers to the host name. Citrix Hypervisor assigns the update file a UUID, which this command prints. Note the UUID.

Tip:

After an update file has been uploaded to the Citrix Hypervisor server, you can use the `update-list` and `update-param-list` commands to view information about the update file.

4. If Citrix Hypervisor detects any errors or preparatory steps that have not been taken, it alerts you. Be sure to follow any guidance before continuing with the update.
5. Update the host, specifying the UUIDs of the host and the update file, by running the following:

```
1 xe update-apply host-uuid=UUID_of_host uuid=UUID_of_file
```

If the host is a member of a pool, ensure that you update the pool master before you update any other pool member.

6. Verify that the update has been successfully applied by using the `update-list` command. If the update has been successful, the `hosts` field contains the host UUID.
7. Perform any post-update operations, as necessary (such as, restarting the XAPI toolstack, or rebooting the host).

Ensure that you apply the update to all hosts in the pool. Running a mixed-mode pool (a pool that includes updated and non-updated servers) is not a supported configuration.

Apply Automated Updates

Automated Updates mode applies any hotfixes and Cumulative Updates that are available for a host. This mode minimizes the number of reboots required to bring the pool or the standalone server pool up-to-date. Where possible, **Automated Updates** mode limits it to a single reboot at the end.

If a new Current Release version is available as an update, **Automated Updates** mode does not apply this update. Instead, you must select manually to update to the new Current Release.

XenCenter requires internet access to fetch the required updates.

To view the list of required updates, perform the following steps:

1. Select the host on the **Resources** pane in XenCenter.
2. Navigate to the **General** tab.
3. Expand the **Updates** section.

You can see:

- **Applied** –lists already-applied updates.
- **Required Updates** –lists the set of updates required to bring the server up-to-date.

Note:

If there are no updates required, the **Required Updates** section is not displayed.

- **Installed supplemental packs** –lists supplemental packs that are installed on the server (if any).

Note:

If you select a pool instead of a server, the **Updates** section lists updates that are already applied as **Fully Applied**.

If you want to choose and install a particular update, see [Apply an update to a pool](#).

Note:

Automated Updates were previously restricted to Citrix Hypervisor Premium Edition customers or Citrix Virtual Apps and Desktops customers. However, in pools with hotfix XS82ECU1053 applied, this feature is available to all users.

Apply Automated Updates by using the Install Update wizard

The following section provides step-by-step instructions on how to apply the set of required updates automatically to bring your pool or standalone host up-to-date.

1. From the XenCenter menu, select **Tools** and then select **Install Update**.
2. Read the information displayed on the **Before You Start** page and then click **Next**.
3. On the **Select Update** page, select the mechanism to use to install the updates. You can see the following options:
 - **Automated Updates** –(default) this option is visible only if XenCenter is connected to at least one licensed pool or a licensed standalone server. Select this option to download and install all the current updates automatically to bring the pool or a standalone server up-to-date.

- **Download update from Citrix** –the Install Update wizard lists available updates from the Support site. To apply the updates, see [Apply an update to a pool](#).
 - **Select update or Supplemental pack from disk** –to install an update you have already downloaded, see [Apply an update to a pool](#). To install supplemental pack updates, see the [Installing Supplemental Packs](#) article in XenCenter documentation.
4. To continue with the automatic application of hotfixes, select **Automated Updates** and then click **Next**.
 5. Select one or more pools or standalone servers that you want to update and click **Next**. Any server or pool that cannot be updated appears unavailable.
 6. The **Install Update** wizard performs several update prechecks, to ensure that the pool is in a valid configuration state.

Follow the on-screen recommendations to resolve any update prechecks that have failed. If you want XenCenter to resolve all failed prechecks automatically, click **Resolve All**. When the prechecks have been resolved, click **Next**.
 7. The Install Update wizard automatically downloads and installs the recommended updates. The wizard also shows the overall progress of the update, displaying the major operations that XenCenter performs while updating each server in the pool.

Notes:

- The updates are uploaded to the default SR of the pool. If the default SR is not shared or does not have enough space, XenCenter tries to upload the update to another shared SR with sufficient space. If none of the shared SRs have sufficient space, the update is uploaded to local storage on each host.
- The update process cannot complete for any reason, XenCenter halts the process. This allows you to fix the issue and resume the update process by clicking the **Retry** button.

8. When all the updates have been applied, click **Finish** to close Install Update wizard.

Live patching in Citrix Hypervisor

The live patching feature applies to hotfixes only. Current Releases and Cumulative Updates cannot be applied as live patches.

Citrix Hypervisor customers who deploy Citrix Hypervisor servers can often be required to reboot their hosts after applying hotfixes. This rebooting results in unwanted downtime for the hosts while customers have to wait until the system is restarted. This unwanted downtime can impact business. Live patching enables customers to install some Linux kernel and Xen hypervisor hotfixes without having

to reboot the hosts. Such hotfixes include both a live patch, which is applied to the memory of the host, and a hotfix that updates the files on disk. Using live patching can reduce maintenance costs and downtime.

When applying an update by using XenCenter, the **Install Update** wizard checks whether the hosts must be rebooted after the update is applied. XenCenter displays the result on the **Prechecks** page. This check enables customers to know the post-update tasks well in advance and schedule the application of hotfixes accordingly.

Note:

Citrix Hypervisor Live Patching is available for Citrix Hypervisor Premium Edition customers, or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. To learn more about Citrix Hypervisor editions, and to find out how to upgrade, visit the [Citrix website](#). For detailed information about Licensing, see [Licensing](#).

Live patching scenarios

Hotfixes can be live patched across pools, hosts, or on a standalone server. Some require a reboot, some require the XAPI toolstack to be restarted, and some hotfixes do not have any post-update tasks. The following scenarios describe the behavior when a Live Patch is and is not available for an update.

- **Updates with a live patch** —Some hotfixes that update the Linux kernel and the Xen hypervisor usually do not require a reboot after applying the hotfix. However, in some rare cases, when the live patch cannot be applied, a reboot might be required.
- **Updates without a live patch** —No change in the behavior here. It works as usual.

Note:

If a host does not require a reboot, or if the hotfix contains live patches, XenCenter displays [No action required](#) on the **Update Mode** page.

Apply Automated Updates and live patching

Automated Updates mode in XenCenter enables you to download and apply the minimum set of hotfixes required to bring your pool or standalone host up-to-date automatically. **Automated Updates** mode does apply any Cumulative Updates that are available for a host. However, if a new Current Release version is available as an update, **Automated Updates** mode does not apply this update. You must manually select to update to the new Current Release.

You can benefit from the live patching feature when you apply hotfixes using the Automated Updates mode in XenCenter. You can avoid rebooting hosts if live patches are available and are successfully applied to the hosts that are updated using **Automated Updates** mode. For more information about the Automated Updates, see [Apply Automated Updates](#).

Enable live patching by using XenCenter and the xe CLI

Live patching feature is enabled by default. Customers can enable or disable live patching using XenCenter or xe CLI command.

Using XenCenter

1. Select the pool or the standalone host on the **Resource** pane.
2. From the **Pool** menu (**Server** in case on standalone hosts) menu, select **Properties** and then click **Live Patching**.
3. On the **Live Patching** page:
 - Select **Use live Patching when possible** to enable live patching.
 - Select **Don't use Live Patching** to disable live patching.

Using the xe CLI

- To enable live patching, run the following command:

```
1 xe pool-param-set live-patching-disabled=false uuid="pool_uuid"
```

- To disable live patching, run the following command:

```
1 xe pool-param-set live-patching-disabled=true uuid="pool_uuid"
```

Troubleshoot the installation

January 9, 2023

Citrix provides two forms of support: free, self-help support from www.citrix.com/support and paid-for Support Services, which you can purchase from the Support site. With Citrix Technical Support, you can open a Support Case online or contact the support center by phone.

The Citrix support site, www.citrix.com/support, hosts various resources. These resources might be helpful to you if you experience odd behavior, crashes, or other problems during installation. Resources include: forums, knowledge base articles, software updates, security bulletins, tools, and product documentation.

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 Citrix Hypervisor Installer
- Press **Alt+F2** to access a local shell
- Press **Alt+F3** to access the event log

If you experience an unknown error during installation, capture the log file from your host and provide it to Technical Support. To capture the log file, complete the following procedure.

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`.

Send the captured log file to the Support team for them to inspect.

Boot from SAN environments

January 9, 2023

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 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.

Boot from SAN depends on SAN-based disk arrays with either hardware Fibre Channel or HBA iSCSI adapter support on the host. For a fully redundant boot from SAN environment, you must configure multiple paths for I/O access. To do so, ensure that the root device has multipath support enabled. For information about whether multipath is available for your SAN environment, consult your storage vendor or administrator. If you have multiple paths available, you can enable multipathing in your Citrix Hypervisor deployment upon installation.

Warning:

Boot-from-SAN settings are *not* inherited during the upgrade process. When upgrading using the ISO or network-boot, follow the same instructions as used in the installation process below to ensure that **multipath** is correctly configured.

To install Citrix Hypervisor to a remote disk on a SAN with multipathing enabled from the installer UI:

1. Boot the computer from the installation media or by using network boot. For more information, see [Install the Citrix Hypervisor server](#)
2. Following the initial boot messages, you see one of the following screens:
 - If you are doing a BIOS installation, you see the Welcome to Citrix Hypervisor screen.



- a) At the welcome screen, press **F2** to select Advanced install.

b) At the boot prompt, enter `multipath`

- If you are doing a UEFI installation, you see a GRUB menu. This menu is shown for 5 seconds.



a) On the GRUB menu, choose `multipath` and press **Enter**.

The Citrix Hypervisor installation process configures the Citrix Hypervisor server, which boots from a remote SAN with multipathing enabled.

To install Citrix Hypervisor to a remote disk on a SAN with multipathing enabled by using a configuration file:

To enable file system multipathing using PXE or UEFI installation, add `device_mapper_multipath=yes` to your configuration file. The following is an example configuration:

```
1 default xenserver
2 label xenserver
3     kernel mboot.c32
4     append /tftpboot/xenserver/xen.gz dom0_max_vcpus=1-2 \
5     dom0_mem=1024M,max:1024M com1=115200,8n1 \
6     console=com1,vga --- /tftpboot/xenserver/vmlinuz \
7     xencons=hvc console=hvc0 console=tty0 \
8     device_mapper_multipath=yes \
9     install --- /tftpboot/xenserver/install.img
```

For additional information on storage multipathing in your Citrix Hypervisor environment, see [Storage](#).

Software-boot-from-iSCSI for Cisco UCS

The Software-boot-from-iSCSI feature enables customers to install and boot Citrix Hypervisor from SAN using iSCSI. Using this feature, Citrix Hypervisor 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.

Citrix Hypervisor supports the following features for Software-boot-from-iSCSI:

- Host installation through PXE-boot

- Cisco UCS vNIC

Software-boot-from-iSCSI has been tested in Legacy BIOS and UEFI boot mode by using Cisco UCS vNICs and Power Vault, NetApp, and EqualLogic arrays. Other configurations might work, however, they have not been validated.

- Jumbo Frames (MTU=9000) configured with the Cisco UCS manager
- Cisco UCS line-rate limiting
- Untagged VLANs
- Networks using the vSwitch back-end
- LVHDoiSCSI SRs and NFS SRs on the same or different SAN/NAS
- Multipathing of the iSCSI root disk
- Compatibility with common Citrix Hypervisor (Network, Maintenance) operations

Requirements

- The primary management interface (IP addressable) and the network for VM traffic, must use separate interfaces.
- Storage (iSCSI targets) must be on a separate Layer 3 (IP) network to all other network interfaces with IP addresses on the host.
- Storage must be on the same subnet as the storage interface of the Citrix Hypervisor server.

Install Citrix Hypervisor by using CD media

Perform the following steps to install Citrix Hypervisor using a CD:

1. Access the boot menu; at the `boot :` prompt enter `menu . c32`
2. Use the cursor keys to select an installation option:

- For a single path LUN, select **install**
- For a multipathed LUN, select **multipath**

3. Press the tab key.

Edit the line ending with the following:

```
1 --- /install.img
```

4. Using the cursor keys, edit this line to read:

```
1 use_ibft --- /install.img
```

5. Press **Enter**.

Citrix Hypervisor server installation proceeds as normal.

Install Citrix Hypervisor by using PXE

Perform the following steps to install Citrix Hypervisor using PXE:

Note:

Ensure that you add the keyword **use_ibft** in the kernel parameters. If multipathing is required, you must add **device_mapper_multipath=enabled**.

The following example shows PXE configuration for a single LUN:

```
1 label xenserver
2 kernel mboot.c32
3 append XS/xen.gz dom0_max_vcpus=2 dom0_mem=1024M,max:1024M
4 com1=115200,8n1 console=com1,vga --- XS/vmlinuz xencons=hvc
   console=tty0
5 console=hvc0 use_ibft --- XS/install.img
```

The following example shows PXE configuration for a multipathed LUN:

```
1 label xenserver
2 kernel mboot.c32
3 append XS/xen.gz dom0_max_vcpus=2 dom0_mem=1024M,max:1024M
4 com1=115200,8n1 console=com1,vga --- XS/vmlinuz xencons=hvc
   console=tty0
5 console=hvc0 use_ibft device_mapper_multipath=enabled --- XS/
   install.img
```

Network boot installations

February 13, 2024

Citrix Hypervisor supports booting hosts using the UEFI mode. UEFI mode provides a rich set of standardized facilities to the bootloader and operating systems. This feature allows Citrix Hypervisor to be more easily installed on hosts where UEFI is the default boot mode.

Note:

- The legacy DOS partition layout is not supported with UEFI boot.
- UEFI Secure Boot is not available for the Citrix Hypervisor host.

The following section contains information about setting up your TFTP and NFS, FTP, or HTTP servers to enable PXE and UEFI booting of Citrix Hypervisor server installations. It then describes how to create an XML answer file, which allows you to perform unattended installations.

Configure your PXE and UEFI environment for Citrix Hypervisor installation

Before you set up the Citrix Hypervisor installation media, configure your TFTP and DHCP servers. The following sections contain information on how to configure your TFTP server for PXE and UEFI booting. Consult your vendor documentation for general setup procedures.

Note:

XenServer 6.0 moved from MBR disk partitioning to GUID Partition Table (GPT). Some third-party PXE deployment systems might attempt to read the partition table on a machine's hard disk *before* deploying the image to the host.

If the deployment system isn't compatible with GPT partitioning scheme and the hard disk has previously been used for a version of Citrix Hypervisor that uses GPT, the PXE deployment system might fail. A workaround for this failure is to delete the partition table on the disk.

In addition to the TFTP and DHCP servers, you require an NFS, FTP, or HTTP server to house the Citrix Hypervisor installation files. These servers can co-exist on one, or be distributed across different servers on the network.

Note:

PXE boot is not supported over a tagged VLAN network. Ensure that the VLAN network you use for PXE boot is untagged.

Additionally, each Citrix Hypervisor server that you want to PXE boot must have a PXE boot-enabled Ethernet card.

The following steps assume that the Linux server you are using has RPM support.

Configure your TFTP server for PXE boot

1. In your TFTP root directory (for example, `/tftpboot`), create a directory called `xenserver`.
2. Copy the `mboot.c32` and `pxelinux.0` files from the `/boot/pxelinux` directory of your installation media to the TFTP root directory.

Note:

We strongly recommend using `mboot.c32` and `pxelinux.0` files from the same source (for example, from the same Citrix Hypervisor ISO).

3. Copy the following files from the Citrix Hypervisor installation media to the new `xenserver` directory on the TFTP server:
 - `install.img` from the root directory
 - `vmlinux` from the `/boot` directory
 - `xen.gz` from the `/boot` directory
4. In the TFTP root directory (for example, `/tftpboot`), create a directory called `pxelinux.cfg`.
5. In the `pxelinux.cfg` directory, create your configuration file called `default`.

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.

Two example configurations are listed below:

- **Example: Unattended install** This example configuration performs an unattended installation using the answer file at the URL specified:

```
1      default xenserver-auto
2      label xenserver-auto
3          kernel mboot.c32
4          append xenserver/xen.gz dom0_max_vcpus=1-16 \
5                  dom0_mem=max:8192M com1=115200,8n1 \
6                  console=com1,vga --- xenserver/vmlinux \
7                  console=hvc0 console=tty0 \
8                  answerfile=<http://pxehost.example.com/
9                      answer_file> \
10                 answerfile_device=<device> \
11                 install --- xenserver/install.img
```

Note:

To specify which network adapter to use to retrieve the answer file, include the `answerfile_device=ethX` or `answerfile_device=MAC` parameter and specify either the Ethernet device number or the MAC address of the device.

For more information about using an answer file, see [Create an answer file for unattended PXE and UEFI installation](#).

- **Example: Manual install** This example configuration starts an installation on any machine that boots from the TFTP server and requires manual responses.

```

1  default xenserver
2  label xenserver
3      kernel mboot.c32
4      append xenserver/xen.gz dom0_max_vcpus=1-16 \
5      dom0_mem=max:8192M com1=115200,8n1 \
6      console=com1,vga --- xenserver/vmlinuz \
7      console=hvc0 console=tty0 \
8      --- xenserver/install.img

```

For more information about PXE configuration file contents, see the [SYSLINUX](#) website.

Configure your TFTP server for UEFI boot

To configure your TFTP server for UEFI boot:

1. In the TFTP root directory (for example, `/tftpboot`), create a directory called `EFI/xenserver`.
2. Configure your DHCP server to provide `/EFI/xenserver/grubx64.efi` as the boot file.
3. Create `grub.cfg` file. For example:

- For an installation that requires manual responses to installation prompts:

```

1  menuentry "Citrix Hypervisor 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/vmlinuz console=hvc0 console=tty0
6      module2 /EFI/xenserver/install.img
7  }

```

- For an unattended installation that uses an answer file:

```

1  menuentry "Citrix Hypervisor 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=eth0 answerfile=http://<ip_address
      >/<path_to_answer_file> install
5     module2 /EFI/xenserver/install.img
6 }
```

For more information about using an answer file, see [Create an answer file for unattended PXE and UEFI installation](#).

4. Copy `grub.cfg` file to `EFI/xenserver` directory on the TFTP server.
5. Copy the following files from the Citrix Hypervisor 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
 - `vmlinuz` from the `/boot` directory
 - `xen.gz` from the `/boot` directory

For details for your specific operating system, see your server operating system manual. The information here is a guide that can be used for Red Hat, Fedora, and some other RPM-based distributions.

To set up the Citrix Hypervisor installation media on an HTTP, FTP or NFS server:

1. On the server, create a directory from which the Citrix Hypervisor installation media can be exported via HTTP, FTP, or NFS.
2. Copy the entire contents of the Citrix Hypervisor installation media to the newly created directory on the HTTP, FTP, or NFS server. This directory is your installation repository.

Note:

When copying the Citrix Hypervisor installation media, ensure that you copy the file `.treeinfo` to the newly created directory.

To prepare the destination system:

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 can proceed unattended.

Install Supplemental Packs during Citrix Hypervisor installation

Supplemental Packs are used to modify and extend the capabilities of Citrix Hypervisor by installing software into the control domain (Dom0). For example, an OEM partner might want to ship Citrix Hypervisor with a set of management tools that require SNMP agents to be installed. Users can add supplemental packs either during initial Citrix Hypervisor installation, or at any time afterwards.

When installing supplemental packs during Citrix Hypervisor installation, unpack each supplemental pack into a separate directory.

Facilities also exist for OEM partners to add their supplemental packs to the Citrix Hypervisor installation repositories to allow automated factory installations.

Create an answer file for unattended PXE and UEFI installation

To perform installations in an unattended fashion, create an XML answer file. 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>sd</guest-disk>
6     <keymap>us</keymap>
7     <root-password>mypassword</root-password>
8     <source type="url">http://pxehost.example.com/citrix-hypervisor
9       /</source>
10    <post-install-script type="url">
11      http://pxehost.example.com/myscripts/post-install-script
12    </post-install-script>
13    <admin-interface name="eth0" proto="dhcp" />
14    <timezone>Europe/London</timezone>
15  </installation>
```

Contain all nodes within a root node named *installation*.

Note:

To enable thin provisioning, specify an `srtype` attribute as `ext`. If this attribute is not specified, the default local storage type is LVM. Thin provisioning sets the local storage type to EXT4 and enables local caching for Citrix Virtual Desktops to work properly. For more information, see [Storage](#).

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   <post-install-script type="url">
6     http://pxehost.example.com/myscripts/post-install-script
7   </post-install-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:

- To enable thin provisioning, specify an `srtype` attribute as `ext`. If this attribute is not specified, the default local storage type is LVM. Thin provisioning sets the local storage type to EXT4 and enables local caching for Citrix Virtual Desktops to work properly. For more information, see [Storage](#).
- 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

<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 Citrix Hypervisor server. If a password is not provided, a prompt is displayed when the server 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 Citrix Hypervisor 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 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, loaded any drivers, 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 `eth0`.
- **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 hostname.

<ntp-server> **Required:** No

Description: Specify one or more NTP servers.

Host partition layout

December 7, 2022

XenServer 7.0 introduced a new host disk partition layout. By moving log files to a larger, separate partition, XenServer can store more detailed logs for a longer time. This feature improves the ability to diagnose issues. Simultaneously, the new partition layout relieves demands on Dom0's root disk and avoids potential space issues due to log file disk space consumption. The default layout contains the following partitions:

- 18 GB Citrix Hypervisor server control domain (dom0) partition
- 18 GB backup partition
- 4 GB logs partition
- 1 GB swap partition
- 0.5 GB UEFI boot partition

In XenServer 6.5 and earlier releases, the 4 GB control domain (dom0) partition was used for all dom0 functions, including swap and logging. Customers who do not use remote syslog or who used with third-party monitoring tools and supplemental packs found the size of the partition to be limited. Citrix Hypervisor eliminates this issue and provides a dedicated 18 GB partition to dom0. In addition, a larger partition dedicated to dom0 reduces demand on the dom0 root disk which can offer significant performance improvements.

The introduction of the 4 GB dedicated log partition eliminates scenarios where excessive logging filled up the dom0 partition and affected the behavior of the host. This partition also enables customers to retain a detailed list of logs for a longer time, improving the ability to diagnose issues.

The partition layout also contains a dedicated 500 MB partition required for UEFI boot.

Note:

If you install Citrix Hypervisor with the new partition layout described above, ensure that you have a disk that is at least 46 GB in size.

To install Citrix Hypervisor on smaller devices, you can do a clean installation of Citrix Hypervisor using the legacy DOS partition layout. A small device is one that has more than 12 GB but less than 46 GB disk space. For more information, see [Installing on Small Devices](#).

Important:

We recommend that you allocate a minimum of 46 GB disk space and install Citrix Hypervisor using the new GPT partition layout.

Legacy partition layout

Note:

The legacy DOS partition layout is deprecated and will be removed in a future release. This partition layout is not supported with UEFI boot.

- XenServer 5.6 Service Pack 2 and earlier used DOS partition tables to separate the root file system and backups from the local storage.
- XenServer 6.0 introduced GUID partition tables to separate root file system, backup and local storage.
- Installing Citrix Hypervisor 8.2 on machines with a required initial partition that must be preserved continues to use the DOS partitioning scheme.

The following table lists the installation and upgrade scenarios and the partition layout that is applied after these operations:

Operation	Number of partitions before upgrade	Number of partitions after installation/upgrade	Partition table type
Clean installation with at least 46 GB of primary disk space	N/A	6	New GPT
Clean installation with <code>disable-gpt</code> with a minimum of 12 GB of primary disk space	N/A	3 (or 4 if there is a utility partition)	DOS
Clean installation on a machine with a utility partition	N/A	3 (or 4 if there is a utility partition)	DOS

Install on small devices

January 10, 2023

Citrix Hypervisor enables customers with smaller devices to install Citrix Hypervisor 8.2 Cumulative Update 1 by using the legacy DOS partition layout. A small device is one that has more than 12 GB but less than 46 GB of disk space. The legacy DOS partition layout includes:

- 4 GB Boot partition
- 4 GB Backup partition
- SR partition (if present on the local disk)

Note:

The legacy DOS partition layout is not supported with UEFI boot.

To install Citrix Hypervisor on small devices, you must add `disable-gpt` to the dom0 parameters. You can use the command `menu . c32` to add the parameter to dom0.

1. Access the boot menu.
2. At the `boot :` prompt, type `menu . c32` and press **Enter**.
3. Use the cursor keys to select an installation option:
 - For a single path LUN, select **install**
 - For a multipathed LUN, select **multipath**
4. Press the **Tab** key.
5. Using the cursor keys, edit the line that ends `/install.img` to include the parameter `disable-gpt` before the last `---`.
For example: `... disable-gpt --- /install.img`
6. Press **Enter**.

Citrix Hypervisor server installation proceeds as normal.

Note:

The installer preserves any utility partition that exists on the host before the installation process.

Important:

We recommend that you allocate a minimum of 46 GB disk space and install Citrix Hypervisor using the new GPT partition layout. For more information, see [Host Partition Layout](#).

Hosts and resource pools

July 15, 2024

This section describes how resource pools can be created through a series of examples using the `xe` command line interface (CLI). A simple NFS-based shared storage configuration is presented and several simple VM management examples are discussed. It also contains procedures for dealing with physical node failures.

Citrix Hypervisor servers and resource pools overview

A *resource pool* comprises multiple Citrix Hypervisor server installations, bound together to a single managed entity which can host Virtual Machines. If combined with shared storage, a resource pool enables VMs to be started on *any* Citrix Hypervisor server which has sufficient memory. The VMs can then be dynamically moved among Citrix Hypervisor servers while running with a minimal downtime (live migration). If an individual Citrix Hypervisor server suffers a hardware failure, the administrator can restart failed VMs on another Citrix Hypervisor server in the same resource pool. When high availability is enabled on the resource pool, VMs automatically move to another host when their host fails. Up to 64 hosts are supported per resource pool, although this restriction is not enforced.

A pool always has at least one physical node, known as the *master*. Only the master node exposes an administration interface (used by XenCenter and the Citrix Hypervisor Command Line Interface, known as the `xe` CLI). The master forwards commands to individual members as necessary.

Note:

When the pool master fails, master re-election takes place only if high availability is enabled.

Requirements for creating resource pools

A resource pool is a homogeneous (or heterogeneous with restrictions) aggregate of one or more Citrix Hypervisor servers, up to a maximum of 64. The definition of homogeneous is:

- CPUs on the server joining the pool are the same (in terms of the vendor, model, and features) as the CPUs on servers already in the pool.
- The server joining the pool is running the same version of Citrix Hypervisor software, at the same patch level, as the servers already in the pool.

The software enforces extra constraints when joining a server to a pool. In particular, Citrix Hypervisor checks that the following conditions are true for the server joining the pool:

- The server is not a member of an existing resource pool.
- The server has no shared storage configured.
- The server is not hosting any running or suspended VMs.
- No active operations are in progress on the VMs on the server, such as a VM shutting down.

- The clock on the server is synchronized to the same time as the pool master (for example, by using NTP).
- The management interface of the server is not bonded. You can configure the management interface when the server successfully joins the pool.
- The management IP address is static, either configured on the server itself or by using an appropriate configuration on your DHCP server.

Citrix Hypervisor servers in resource pools can contain different numbers of physical network interfaces and have local storage repositories of varying size. In practice, it is often difficult to obtain multiple servers with the exact same CPUs, and so minor variations are permitted. If it is acceptable to have hosts with varying CPUs as part of the same pool, you can force the pool-joining operation by passing the `--force` parameter.

All hosts in the pool must be in the same site and connected by a low latency network.

Note:

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

A pool must contain shared storage repositories to select on which Citrix Hypervisor server to run a VM and to move a VM between Citrix Hypervisor servers dynamically. If possible create a pool after shared storage is available. We recommend that you move existing VMs with disks located in local storage to shared storage after adding shared storage. You can use the `xe vm-copy` command or use XenCenter to move VMs.

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 is 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](#).

Note:

You can only join a new host to a resource pool when the host's management interface is on the same tagged VLAN as that of the resource pool.

Add a host to a pool by using the xe CLI

1. Open a console on the Citrix Hypervisor host that you want to join to a pool.
2. Join the Citrix Hypervisor host to the pool by issuing the command:

```
1 xe pool-join master-address=<address of pool master> master-  
   username=<administrator username> master-password=<password>
```

The `master-address` must be set to the fully qualified domain name of the pool master. The `password` must be the administrator password set when the pool master 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 master.

Citrix Hypervisor servers 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

Citrix Hypervisor 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 Citrix Hypervisor servers 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 servers are not required to be the same type at the level of family, model, or stepping numbers.

Citrix Hypervisor 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, Citrix Hypervisor 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

For a complete list of supported shared storage types, see [Storage repository formats](#). This section shows how shared storage (represented as a storage repository) can be created on an existing NFS server.

To add NFS shared storage to a resource pool by using the CLI

1. Open a console on any Citrix Hypervisor server 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 Citrix Hypervisor server in the pool. Any Citrix Hypervisor servers 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 [Storage repository formats](#).

Remove Citrix Hypervisor servers from a resource pool

Note:

Before removing any Citrix Hypervisor server 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 Citrix Hypervisor servers 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 Citrix Hypervisor server 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 Citrix Hypervisor servers 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 Citrix Hypervisor servers. The VMs do not start until the virtual disks associated with them have been changed to point at shared storage seen by other Citrix Hypervisor servers 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 Citrix Hypervisor servers 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.

Prepare a pool of Citrix Hypervisor servers 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 Citrix Hypervisor server in the pool. You can do this by placing the Citrix Hypervisor server 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 master host in to maintenance mode results in the loss of the last 24 hrs of RRD updates for offline VMs.

Warning:

We highly recommend rebooting all Citrix Hypervisor servers before installing an update and then verifying their configuration. Some configuration changes only take effect when the Citrix Hypervisor server is rebooted, so the reboot may 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=Citrix Hypervisor_host_uuid
2 xe host-evacuate uuid=Citrix Hypervisor_host_uuid
```

This command disables the Citrix Hypervisor server and then migrates any running VMs to other Citrix Hypervisor servers in the pool.

2. Perform the desired maintenance operation.
3. Enable the Citrix Hypervisor server when the maintenance operation is complete:

```
1 xe host-enable
```

4. Restart any halted VMs and resume any suspended VMs.

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, servers, 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 Citrix Hypervisor Premium Edition customers, or those who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.

The list of resources and various types of resource data that are included in the report:

Server:

- Name
- Pool Master
- 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 Citrix Hypervisor server.

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**.

Host power-on

Powering on hosts remotely

You can use the Citrix Hypervisor server Power On feature to turn a server 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.**
- **Dell Remote Access Cards (DRAC).** To use Citrix Hypervisor with DRAC, you must install the Dell supplemental pack to get DRAC support. DRAC support requires installing the RACADM command-line utility on the server with the remote access controller and enabling DRAC and its interface. RACADM is often included in the DRAC management software. For more information, see Dell's DRAC documentation.
- A custom script based on the management API that enables you to turn the power on and off through Citrix Hypervisor. 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 a DRAC card, 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 Citrix Hypervisor).

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", "DRAC","custom") \  
3     power-on-config=key:value
```

For DRAC the keys are `power_on_ip` to specify the password if you are using the secret feature. For more information, see [Secrets](#).

To turn on hosts remotely by using the CLI Run the command:

```
1 xe host-power-on host=<host uuid>
```

Configure a custom script for the Host Power On feature

If your server's remote-power solution uses a protocol that is not supported by default (such as Wake-On-Ring or Intel Active Management Technology), you can create a custom Linux Python script to turn on your Citrix Hypervisor computers remotely. However, you can also create custom scripts for DRAC and Wake on LAN remote-power solutions.

This section provides information about configuring a custom script for Host Power On using the key/-value pairs associated with the Citrix Hypervisor API call `host.power_on`.

When you create a custom script, run it from the command line each time you want to control power remotely on a Citrix Hypervisor server. Alternatively, you can specify it in XenCenter and use the XenCenter UI features to interact with it.

The Citrix Hypervisor API is documented in the document, the Citrix Hypervisor Management API, which is available from the [developer documentation](#) website.

Warning:

Do not change the scripts provided by default in the `/etc/xapi.d/plugins/` directory. You can include new scripts in this directory, but you must never change the scripts contained in that directory after installation.

Key/Value Pairs To use Host Power On, configure the `host.power_on_mode` and `host.power_on_config` keys. See the following section for information about the values.

There is also an API call that lets you set these fields simultaneously:

```
1 void host.set_host_power_on_mode(string mode, Dictionary<string,string> config)
```

host.power_on_mode

- **Definition:** Contains key/value pairs to specify the type of remote-power solution (for example, Dell DRAC).
- **Possible values:**
 - An empty string, representing power-control disabled.
 - “DRAC”: Lets you specify Dell DRAC. To use DRAC, you must have already installed the Dell supplemental pack.
 - “wake-on-lan”: Lets you specify Wake on LAN.
 - Any other name (used to specify a custom power-on script). This option is used to specify a custom script for power management.
- **Type:** string

host.power_on_config

- **Definition:** Contains key/value pairs for mode configuration. Provides additional information for DRAC.
- **Possible values:**
 - If you configured DRAC as the type of remote-power solution, you must also specify one of the following keys:
 - ★ “power_on_ip”: The IP address you specified configured to communicate with the power-control card. Alternatively, you can type the domain name for the network interface where DRAC is configured.

- “power_on_user”: The DRAC user name associated with the management processor, which you may have changed from its factory default settings.
- “power_on_password_secret”: Specifies using the secrets feature to secure your password.
 - To use the secrets feature to store your password, specify the key “power_on_password_secret”. For more information, see [Secrets](#).
- **Type:** Map (string, string)

Sample script The sample script imports the Citrix Hypervisor API, defines itself as a custom script, and then passes parameters specific to the host you want to control remotely. You must define the parameters `session` in all custom scripts.

The result appears when the script is unsuccessful.

```
1 import XenAPI
2 def custom(session,remote_host,
3 power_on_config):
4 result="Power On Not Successful"
5 for key in power_on_config.keys():
6 result=result+' '
7 key=' '+key+' '
8 value=' '+power_on_config[key]
9 return result
```

Note:

After creating the script, save it in `/etc/xapi.d/plugins` with a `.py` extension.

Communicate with Citrix Hypervisor servers and resource pools

TLS

Citrix Hypervisor uses the TLS 1.2 protocol to encrypt management API traffic. Any communication between Citrix Hypervisor 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.

Citrix Hypervisor uses the following cipher suite:

- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

In addition, the following cipher suites are also supported for backwards compatibility with some versions of Citrix Virtual Apps and Desktops:

- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA256

Note:

These additional cipher suites use CBC mode. Although some organizations prefer GCM mode, Windows Server 2012 R2 does not support RSA cipher suites with GCM mode. Clients running on Windows Server 2012 R2 that connect to a Citrix Hypervisor server or pool, might need to use these CBC-mode cipher suites.

SSH

When using an SSH client to connect directly to the Citrix Hypervisor server the following algorithms can be used:

Ciphers:

- aes128-ctr
- aes256-ctr
- aes128-gcm@openssh.com
- aes256-gcm@openssh.com
- aes128-cbc
- aes256-cbc

MACs:

- hmac-sha2-256
- hmac-sha2-512
- hmac-sha1

KexAlgorithms:

- curve25519-sha256
- ecdh-sha2-nistp256
- ecdh-sha2-nistp384
- ecdh-sha2-nistp521
- diffie-hellman-group14-sha1

HostKeyAlgorithms:

- ecdsa-sha2-nistp256
- ecdsa-sha2-nistp384
- ecdsa-sha2-nistp521

- ssh-ed25519
- ssh-rsa

If you want to disable SSH access to your Citrix Hypervisor server, you can do this in [xsconsole](#).

1. From XenCenter, open the server console and log in as [root](#).
2. Type [xsconsole](#).
3. In [xsconsole](#), go to **Remote Service Configuration > Enable/Disable Remote Shell**.

The console displays whether remote shell is enabled.

4. To change whether the remote shell is enabled or disabled, press **Enter**.

Important:

We do not support customer modifications to the cryptographic functionality of the product.

Install a TLS certificate on your server

The Citrix Hypervisor server comes installed with a default TLS certificate. However, to use HTTPS to secure communication between Citrix Hypervisor 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 passkey.
- The key and certificate are in base-64 encoded 'PEM' format.
- The certificate is valid and has not expired.
- The signature algorithm is SHA-2 (SHA256).

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 Citrix Hypervisor server.
- 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 Citrix Hypervisor server 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 Citrix Hypervisor server

1. Generate a certificate signing request First, generate a private key and certificate signing request. On the Citrix Hypervisor server, complete the following steps:

1. To create a private key file, run the following command:

```
1 openssl genrsa -des3 -out privatekey.pem 2048
```

2. Remove the password 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 Citrix Hypervisor server. Citrix recommends 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 Fu40qMSm1dNCFfnACRwYQkQgqCt/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 is a trusted third-party that provides digital certificates. 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
- a root certificate
- if applicable, an intermediate certificate

You can now install all these files on your Citrix Hypervisor server.

3. Install the signed certificate on your Citrix Hypervisor server After the certificate authority responds to the certificate signing request, complete the following steps to install the certificate on your Citrix Hypervisor server:

1. Get the signed certificate, root certificate and, if the certificate authority has one, the intermediate certificate from the certificate authority.
2. Copy the key and certificates to the Citrix Hypervisor server.
3. Run the following command on the server:

```
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 additional security, you can delete the private key file after the certificate is installed.

Manage the administrator password

When you first install a Citrix Hypervisor server, you set an administrator or *root* password. You use this password to connect XenCenter to your server or (with user name `root`) to log into **xsconsole**, the system configuration console.

If you join a server to a pool, the administrator password for the server is automatically changed to match the administrator password of the pool master.

Note:

Citrix Hypervisor 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 server by using XenCenter, complete the following steps:

1. In the **Resources** pane, select the pool or any server 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 server, select the server in the **Resources** pane, and click **Password** and then **Change** from the **Server** menu.

If XenCenter is configured to save your server login credentials between sessions, the new password is remembered. For more information, see [Store your server 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 server 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 server by using **xsconsole**, complete the following steps:

1. On the pool master, 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 server is pool master, this updated password is now propagated to the other servers 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 Citrix Hypervisor server, you can reset the password by accessing the server directly.

1. Reboot the Citrix Hypervisor server.
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 server is pool master, this updated password is now propagated to the other servers in the pool.

After changing the administrator password, rotate the pool secret.

Rotate the pool secret

The pool secret is a secret shared among the servers in a pool that enables the server 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 server 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. 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 server 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](#).

Enable IGMP snooping on your Citrix Hypervisor pool

Citrix Hypervisor 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:

- IGMP snooping is available only when the network back-end uses Open vSwitch.
- 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 Citrix Hypervisor 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 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 was 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 should 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.

Clustered pools

May 20, 2024

Clustering provides extra features that are required for resource pools that use GFS2 SRs. For more information about GFS2, see [Configure storage](#).

A cluster is a pool of Citrix Hypervisor 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.

Note:

The clustering feature only benefits pools that contain a GFS2 SR. If your pool does not contain a GFS2 SR, do not enable clustering in your pool.

Quorum

Each host in a cluster must always be in communication with the majority of hosts in the 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 cluster requires to keep quorum.

The following table summarizes this behavior. The value of *n* is the total number of hosts in the 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	$(n/2)+1$	$n/2$

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 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 clustered pool, 3 hosts must be contactable for the 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 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 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 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 cluster with an even number of hosts to split exactly in half. The running cluster decides which half of the cluster self-fences and which half of the cluster has quorum. The half of the cluster that contains the node with the lowest ID that was seen as active before the cluster split remains active and the other half of the cluster self-fences.

For example, in a 4-host clustered pool, 3 hosts must be contactable for the cluster to become active $[4/2 + 1 = 3]$. After the 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 clustered pool that uses high availability, Citrix Hypervisor restarts the VMs according to their restart configuration on other pool members. The host that self-fenced restarts and attempts to rejoin the cluster.

If the number of live hosts in the cluster becomes less than the quorum value, all the remaining hosts lose quorum.

In an ideal scenario, your clustered pool always has more live hosts than are required for quorum and Citrix Hypervisor never fences. To make this scenario more likely, consider the following recommendations when setting up your clustered pool:

- Ensure that you have good hardware redundancy.

- Use a dedicated bonded network for the 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 clustered pool

Before you begin, ensure the following prerequisites are met:

- All Citrix Hypervisor servers in the 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 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 cluster must use static IP addresses for the cluster network.
- 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.
- 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 Citrix technologies](#).

- If you are clustering an existing pool, ensure that high availability is disabled. You can enable high availability again after clustering is enabled.
- We strongly recommend that you use a bonded network for your clustered pool that is not used for any other traffic.

If you prefer, you can set up clustering on your pool by using XenCenter. For more information, see the [XenCenter product documentation](#).

To use the xe CLI to create a clustered pool:

1. Create a bonded network to use as the clustering network.

Note:

We strongly recommend that you use a dedicated bonded network for your clustered pool. Do not use this network for any other traffic.

On the Citrix Hypervisor server that you want to be the pool master, complete the following steps:

- a) Open a console on the Citrix Hypervisor server.
- b) Name your resource pool by using the following command:

```
1 xe pool-param-set name-label="New Pool" uuid=<pool_uuid>
```

- c) 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.

- d) Find the UUIDs of the PIFs to use in the bond by using the following command:

```
1 xe pif-list
```

- e) 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 master, when you join other Citrix Hypervisor servers 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 Citrix Hypervisor servers.

Repeat the following steps on each Citrix Hypervisor server that is a (non-master) pool member:

- a) Open a console on the Citrix Hypervisor server.
- b) Join the Citrix Hypervisor server to the pool on the pool master 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 Citrix Hypervisor server that is the pool master. The `password` must be the administrator password set when the pool master 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 Citrix Hypervisor server in your resource pool:

```
1 xe pif-param-set disallow-unplug=true uuid=<pif_uuid>
```

4. Enable clustering on your pool. Run the following command on a Citrix Hypervisor server 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.

Destroy a clustered pool

You can destroy a clustered pool. After you destroy a clustered pool, the pool continues to exist, but is no longer clustered and can no longer use GFS2 SRs.

To destroy a clustered pool, run the following command:

```
1 xe cluster-pool-destroy cluster-uuid=<uuid>
```

Manage your clustered pool

When managing your clustered pool, the following practices can decrease the risk of the pool losing quorum.

Add or remove a host on a clustered pool

When adding or removing a host on a clustered pool, ensure that all the hosts in the cluster are on-line.

You can add or remove a host on a 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 clustered pool by using the xe CLI. For more information, see [Add a host to a pool by using the xe CLI](#) and [Remove Citrix Hypervisor 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 cluster until it is started again. While the host is shut down, it does not count toward the quorum value of the cluster. The host absence does not cause other hosts to lose quorum.

However, if a host is forcibly or unexpectedly shut down, it is not removed from the cluster before it goes offline. This host does count toward the quorum value of the 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 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 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 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 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 master, 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 cluster.

You can only change the IP address of a host that is part of a clustered pool when that host is in maintenance mode. Changing the IP address of a host causes the host to leave the cluster. When the IP address has been successfully changed, the host rejoins the cluster. After the host rejoins the 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 cluster members are offline, they count towards the quorum number for the cluster and decrease the number of cluster members that are contactable. This situation increases the risk of a subsequent host failure causing the cluster to lose quorum and shut down completely.

Having offline hosts in your cluster also prevents you from performing certain actions. In a clustered pool, every member of the pool must agree to every change of pool membership before the change can be successful. If a cluster member is not contactable, Citrix Hypervisor prevents operations that change 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 clustered pool to forget them. These hosts are permanently removed from the pool. After hosts are removed from the 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 clustered pool is told to forget a host, the host cannot be added back into the pool.

To rejoin the 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 clustered pool in the usual way.

Troubleshoot your clustered pool

If you encounter issues with your clustered pool, see [Troubleshoot clustered pools](#).

Constraints

- Clustered pools only support up to 16 hosts per pool.
- To enable HA in a clustered pool, the heartbeat SR must be a GFS2 SR.
- For cluster traffic, you must 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 clustering network while the cluster is live and has running VMs. This action can cause the cluster to fence.
- If you have an IP address conflict (multiple hosts having the same IP address) on your clustering network involving at least one host with clustering enabled, the hosts do not fence. To fix this issue, resolve the IP address conflict.

Troubleshoot clustered pools

June 19, 2024

Citrix Hypervisor 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 Citrix Hypervisor 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 no host in the pool is offline when a new host is trying to join the clustered pool.
5. 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-plugin 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 marked as dead, it can't be added back into the cluster. To add this host back into the cluster, you must do a fresh installation of Citrix Hypervisor on the host.

I've repaired a host that was marked as dead. How do I add it back into my cluster?

A Citrix Hypervisor 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 Citrix Hypervisor hosts in the cluster gets into a fence loop because of continuously losing and gaining quorum, you can boot the host with the `nocluster` kernel command-line argument. Connect to the physical or serial console of the host and edit the boot arguments in grub.

Example:

```

1 /boot/grub/grub.cfg
2 menuentry 'XenServer' {
3
4     search --label --set root root-oyftuj
5     multiboot2 /boot/xen.gz dom0_mem=4096M,max:4096M watchdog ucode
      =scan dom0_max_vcpus=1-16 crashkernel=192M,below=4G console=
      vga vga=mode-0x0311
6     module2 /boot/vmlinuz-4.4-xen root=LABEL=root-oyftuj ro nolvm
      hpet=disable xencons=hvc console=hvc0 console=tty0 quiet vga
      =785 splash plymouth.ignore-serial-consoles nocluster
7     module2 /boot/initrd-4.4-xen.img
8 }
9
10 menuentry 'Citrix Hypervisor (Serial)' {
11
12     search --label --set root root-oyftuj
13     multiboot2 /boot/xen.gz com1=115200,8n1 console=com1,vga
      dom0_mem=4096M,max:4096M watchdog ucode=scan dom0_max_vcpus
      =1-16 crashkernel=192M,below=4G
14     module2 /boot/vmlinuz-4.4-xen root=LABEL=root-oyftuj ro nolvm
      hpet=disable console=tty0 xencons=hvc console=hvc0 nocluster
15     module2 /boot/initrd-4.4-xen.img
16 }

```

What happens when the pool master gets restarted in a clustered pool?

In most cases, the behavior when the pool master 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 master is selected and general service is maintained.
- If HA is not enabled, there is no master for the pool. Running VMs on the remaining hosts continue to run. Most administrative operations aren't available until the master 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 [Citrix Hypervisor 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.

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 reenabling 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

February 22, 2024

Defining users, groups, roles and permissions allows you to control who has access to your Citrix Hypervisor servers and pools and what actions they can perform.

When you first install Citrix Hypervisor, a user account is added to Citrix Hypervisor automatically. This account is the local super user (LSU), or root, which Citrix Hypervisor authenticates locally.

The LSU, or root, is a special user account intended for system administration and has all permissions. In Citrix Hypervisor, the LSU is the default account at installation. Citrix Hypervisor 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 Citrix Hypervisor 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, Citrix Hypervisor does not assign newly created user accounts RBAC roles automatically. Therefore, these accounts do not have any access to the Citrix Hypervisor pool until you assign them a role.

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.

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 server or a pool, you must use Active Directory user accounts for authentication. AD accounts let Citrix Hypervisor 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](#).

You can configure varying levels of access for specific users by enabling Active Directory authentication, adding user accounts, and assign 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 Citrix Hypervisor maps to an entity on your directory server (either a user or a group). When you enable external authentication, Citrix Hypervisor 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 Citrix Hypervisor CLI uses slightly different terminology to refer to Active Directory and user account features: XenCenter Term Citrix Hypervisor CLI Term
Users Subjects Add users Add subjects

Even though Citrix Hypervisor is Linux-based, Citrix Hypervisor lets you use Active Directory accounts for Citrix Hypervisor user accounts. To do so, it passes Active Directory credentials to the Active Directory domain controller.

When you add Active Directory to Citrix Hypervisor, Active Directory users and groups become Citrix Hypervisor 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 Citrix Hypervisor. Users and groups do not need to qualify their user name by using a domain name.

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 Citrix Hypervisor) first.

The external authentication process works as follows:

1. The credentials supplied when connecting to a server 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 Citrix Hypervisor, 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

Citrix Hypervisor supports use of Active Directory servers using Windows 2008 or later.

To authenticate Active Directory for Citrix Hypervisor servers, you must use the same DNS server for both the Active Directory server (configured to allow interoperability) and the Citrix Hypervisor server. 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 Citrix Hypervisor server. 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:

Citrix Hypervisor server names must be unique throughout the Citrix Hypervisor deployment.

Note the following:

- Citrix Hypervisor labels its AD entry on the AD database using its hostname. If two Citrix Hypervisor servers with the same hostname are joined to the same AD domain, the second Citrix Hypervisor overwrites the AD entry of the first Citrix Hypervisor. The overwriting occurs regardless of whether the hosts belong to the same or different pools. This can cause the AD authentication on the first Citrix Hypervisor to stop working.

You can use the same host name in two Citrix Hypervisor servers, as long as they join different AD domains.

- The Citrix Hypervisor servers 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 Citrix Hypervisor pool and the Active Directory server.
- Mixed-authentication pools are not supported. You cannot have a pool where some servers in the pool are configured to use Active Directory and some are not).
- The Citrix Hypervisor Active Directory integration uses the Kerberos protocol to communicate with the Active Directory servers. Therefore, Citrix Hypervisor 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 Citrix Hypervisor servers must be synchronized with the clocks on your Active Directory server. When Citrix Hypervisor 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 consist solely of no more than 63 alphanumeric characters, and must not be purely numeric.

When you add a server to a pool after enabling Active Directory authentication, you are prompted to configure Active Directory on the server joining the pool. When prompted for credentials on the joining server, type Active Directory credentials with sufficient privileges to add servers to that domain.

Active Directory integration

Ensure that the following firewall ports are open for outbound traffic in order for Citrix Hypervisor to access the domain controllers.

Port	Protocol	Use
53	UDP/TCP	DNS
88	UDP/TCP	Kerberos 5
123	UDP	NTP
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

For more information, see [Communication Ports Used by Citrix Hypervisor](#).

Notes:

- To view the firewall rules on a Linux computer using *iptables*, run the following command:
`iptables -nL`.
- Citrix Hypervisor uses PowerBroker Identity Services (PBIS) to authenticate the AD user in the AD server, and to encrypt communications with the AD server.

How does Citrix Hypervisor manage the machine account password for AD integration?

Similarly to Windows client machines, PBIS automatically updates the machine account password. PBIS renews the password every 30 days, or as specified in the machine account password renewal policy in the AD server.

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 Citrix Hypervisor servers, 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 Citrix Hypervisor servers:

```
1 xe pif-reconfigure-ip mode=static dns=dns host 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 Citrix Hypervisor to deal with failures that occur when enabling authentication on a particular host. Citrix Hypervisor 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
```

User authentication

To allow a user access to your Citrix Hypervisor server, 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 Citrix Hypervisor, 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 Citrix Hypervisor:

```
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, `xendt\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 Citrix Hypervisor server or pool:

```
1 xe subject-list
```

Remove access for a user

When a user is authenticated, they can access the server 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 server 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

January 9, 2023

The Role-Based Access Control (RBAC) feature in Citrix Hypervisor allows you to assign users, roles, and permissions to control who has access to your Citrix Hypervisor and what actions they can perform. The Citrix Hypervisor RBAC system maps a user (or a group of users) to defined roles (a named set of permissions). The roles have associated Citrix Hypervisor 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. Citrix Hypervisor 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, Citrix Hypervisor 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 Citrix Hypervisor. The LSU is authenticated through Citrix Hypervisor 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 Citrix Hypervisor 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

May 16, 2024

Roles

Citrix Hypervisor 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 server 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.

Warning:

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 Citrix Hypervisor, you must first assign a role to the newly created user before they can use the account. Citrix Hypervisor **does not** automatically assign a role to the newly created user. As a result, these accounts do not have any access to Citrix Hypervisor 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/modify roles						
Log in to (physical) server consoles (through SSH and XenCenter)	X					
Server backup/re-store	X					

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Import OVF/OVA packages; import disk images	X					
Import XVA packages	X	X	X			
Export OVF/OVA packages	X					
Export XVA packages	X	X	X			
Set cores per socket	X	X	X	X		
Convert VMs using XenServer Conversion Manager	X					
Switch-port locking	X	X				
Multipathing	X	X				
Log out active user connec- tions	X	X				
Create and dismiss alerts	X	X				
Cancel task of any user	X	X				
Pool man- agement	X	X				
Live migration	X	X	X			

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
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	
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
Connect to pool and read all pool metadata	X	X	X	X	X	X
Configure virtual GPU	X	X				
View virtual GPU config- uration	X	X	X	X	X	X
Access the config drive (CoreOS VMs only)	X					

Role per- missions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Scheduled Snapshots (Add/Re- move VMs to existing Snapshots Schedules)	X	X	X			
Scheduled Snapshots (Add/Modi- fy/Delete Snapshot Schedules)	X	X				
Gather diagnostic information	X	X				
Configure changed block tracking	X	X	X	X		
List changed blocks	X	X	X	X	X	
Configure PVS- Accelerator	X	X				
View PVS- Accelerator configura- tion	X	X	X	X	X	X

Definitions of permissions

Assign/modify roles:

- Add/remove users
- Add/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 VM create/destroy operations:

- Back up and restore servers
- Back up and restore pool metadata

The capability to restore a backup lets the assignee revert RBAC configuration changes.

Import/export OVF/OVA packages and disk images:

- Import OVF and OVA packages
- Import disk images
- Export VMs as OVF/OVA 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 XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager):

- Convert VMware ESXi/vCenter VMs to Citrix Hypervisor VMs

This permission lets the user convert workloads from VMware to Citrix Hypervisor by copying batches of VMware ESXi/vCenter VMs to Citrix Hypervisor 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 is allowed to send traffic.

Multipathing:

- Enable multipathing
- Disable multipathing

Log out active user connections:

- Ability to disconnect logged in users

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 Citrix Hypervisor cancel an in-progress task initiated by any user.

Pool management:

- Set pool properties (naming, default SRs)
- Create a clustered pool
- Enable, disable, and configure high availability
- Set per-VM high availability 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 master
- Emergency master address
- Emergency recover pool members
- Designate new master
- 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, because of maintenance mode, or high availability
- Configure server management interface 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 server if they are not satisfied with the server Citrix Hypervisor selected.

VM create/destroy operations:

- Install or delete
- Clone/copy VMs
- Add, remove, and configure virtual disk/CD devices
- Add, remove, and configure virtual network devices
- Import/export XVA files
- VM configuration change
- Server backup/restore

VM change CD media:

- Eject current CD
- Insert new CD

Import/export OVF/OVA packages; import disk images

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 server consoles.

XenCenter view management operations:

- Create and modify global XenCenter folders
- Create and modify global XenCenter custom fields
- Create and modify global XenCenter searches

Folders, custom fields, and searches are shared between all users accessing the pool

Cancel own tasks:

- Lets a user cancel their own tasks

Read audit log:

- Download the Citrix Hypervisor audit log

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 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

Access the config drive (CoreOS VMs only):

- Access the config driver of the VM

Scheduled Snapshots:

- Add VMs to existing snapshot schedules
- Remove VMs from existing snapshot schedules
- Add snapshot schedules
- Modify snapshot schedules
- Delete snapshot schedules

Gather diagnostic information from Citrix Hypervisor:

- 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

Changed block tracking can be enabled only for licensed instances of Citrix Hypervisor 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/Remove (PVS-Accelerator) cache configuration

View PVS-Accelerator configuration:

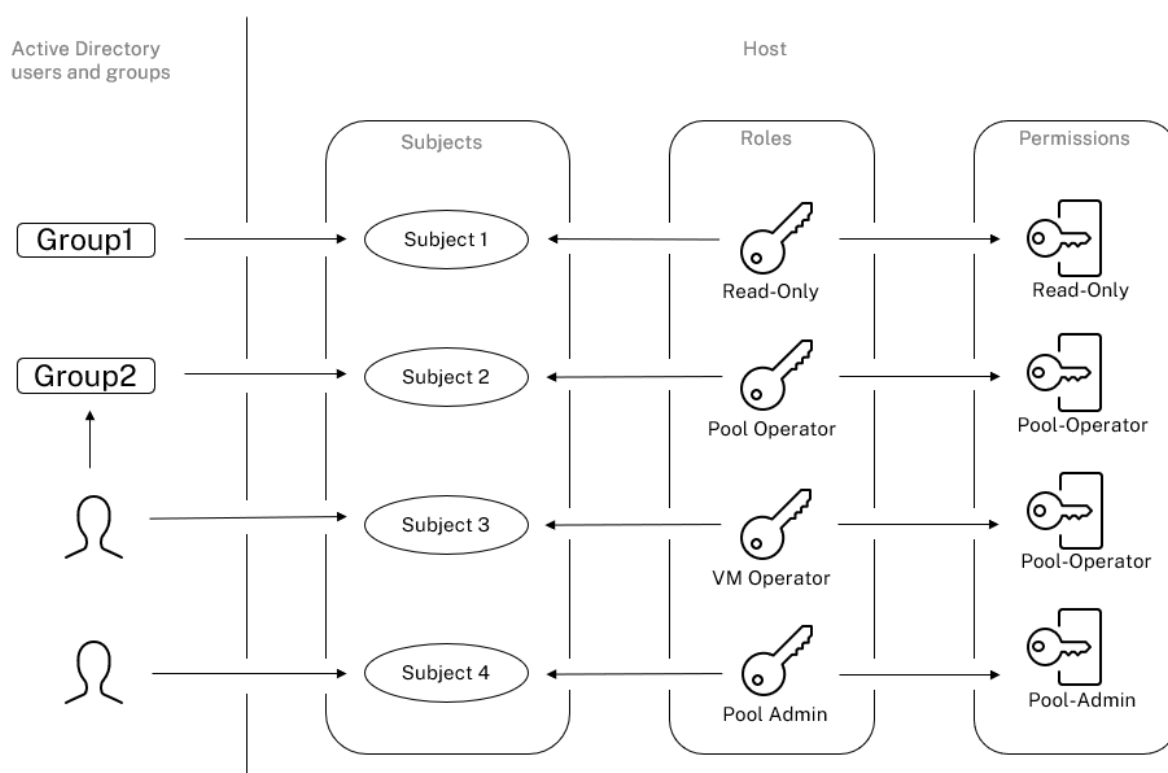
- View the status of PVS-Accelerator

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 Citrix Hypervisor 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. Citrix Hypervisor 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

January 9, 2023

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( RO): 0165f154-ba3e-034e-6b27-5d271af109ba
2  name ( RO): pool-admin
3  description ( RO): The Pool Administrator role has full access to
4  all
   features and settings, including accessing Dom0 and managing
   subjects,
```

```
5     roles and external authentication
6
7     uuid ( R0): b9ce9791-0604-50cd-0649-09b3284c7dfd
8     name ( R0): pool-operator
9     description ( R0): The Pool Operator role manages host- and pool-
10        wide resources,
11        including setting up storage, creating resource pools and managing
12        patches, and
13        high availability (HA).
14
15     uuid( R0): 7955168d-7bec-10ed-105f-c6a7e6e63249
16     name ( R0): vm-power-admin
17     description ( R0): The VM Power Administrator role has full access
18        to VM and
19        template management and can choose where to start VMs and use the
20        dynamic memory
21        control and VM snapshot features
22
23     uuid ( R0): aaa00ab5-7340-bfbc-0d1b-7cf342639a6e
24     name ( R0): vm-admin
25     description ( R0): The VM Administrator role can manage VMs and
26        templates
27
28     uuid ( R0): fb8d4ff9-310c-a959-0613-54101535d3d5
29     name ( R0): vm-operator
30     description ( R0): The VM Operator role can use VMs and interact
31        with VM consoles
32
33     uuid ( R0): 7233b8e3-eacb-d7da-2c95-f2e581cdbf4e
34     name ( R0): read-only
35     description ( R0): The Read-Only role can log in with basic read-
36        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 Citrix Hypervisor users, their uuid, and the roles they are associated with:

```
1     uuid ( R0): bb6dd239-1fa9-a06b-a497-3be28b8dca44
2     subject-identifier ( R0): S
3     -1-5-21-1539997073-1618981536-2562117463-2244
4     other-config (MRO): subject-name: example01\user_vm_admin; subject-
5        upn: \
```

```

4      user_vm_admin@XENDT.NET; subject-uid: 1823475908; subject-gid:
5      1823474177; \
6      subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2244;
7      subject-gecos: \
8      user_vm_admin; subject-displayname: user_vm_admin; subject-is-
9      group: false; \
10     subject-account-disabled: false; subject-account-expired: false;
11     \
12     subject-account-locked: false; subject-password-expired: false
13     roles (SRO): vm-admin
14
15     uuid ( RO): 4fe89a50-6a1a-d9dd-afb9-b554cd00c01a
16     subject-identifier ( RO): S
17     -1-5-21-1539997073-1618981536-2562117463-2245
18     other-config (MRO): subject-name: example02\user_vm_op; subject-upn
19     : \
20     user_vm_op@XENDT.NET; subject-uid: 1823475909; subject-gid:
21     1823474177; \
22     subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2245; \
23     subject-gecos: user_vm_op; subject-displayname: user_vm_op; \
24     subject-is-group: false; subject-account-disabled: false; \
25     subject-account-expired: false; subject-account-locked: \
26     false; subject-password-expired: false
27     roles (SRO): vm-operator
28
29     uuid ( RO): 8a63fbf0-9ef4-4fef-b4a5-b42984c27267
30     subject-identifier ( RO): S
31     -1-5-21-1539997073-1618981536-2562117463-2242
32     other-config (MRO): subject-name: example03\user_pool_op; \
33     subject-upn: user_pool_op@XENDT.NET; subject-uid: 1823475906; \
34     subject-gid: 1823474177; subject-s id:
35     S-1-5-21-1539997073-1618981536-2562117463-2242; \
36     subject-gecos: user_pool_op; subject-displayname: user_pool_op; \
37     subject-is-group: false; subject-account-disabled: false; \
38     subject-account-expired: false; subject-account-locked: \
39     false; subject-password-expired: false
40     roles (SRO): pool-operator

```

To add a subject to RBAC

To enable existing AD users to use RBAC, create a subject instance within Citrix Hypervisor, 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 “Log-out 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 server 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
```

Networking

April 25, 2024

This section provides an overview of Citrix Hypervisor networking, including networks, VLANs, and NIC bonds. It also discusses how to manage your networking configuration and troubleshoot it.

Important:

vSwitch is the default network stack of Citrix Hypervisor. Follow the instructions in [vSwitch networks](#) to configure the Linux network stack.

If you are already familiar with Citrix Hypervisor networking concepts, you can skip ahead to [Manage networking](#) for information about the following sections:

- Create networks for standalone Citrix Hypervisor servers
- Create networks for Citrix Hypervisor servers that are configured in a resource pool
- Create VLANs for Citrix Hypervisor servers, either standalone or part of a resource pool
- Create bonds for standalone Citrix Hypervisor servers
- Create bonds for Citrix Hypervisor servers 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

Citrix Hypervisor 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:

Citrix Hypervisor provides automated configuration and management of NICs using the xe command line interface (CLI). Do not edit the host networking configuration files directly.

vSwitch networks

vSwitch networks support open flow.

- Supports fine-grained security policies to control the flow of traffic sent to and from a VM.
- Provides detailed visibility about the behavior and performance of all traffic sent in the virtual network environment.

A vSwitch greatly simplifies IT administration in virtualized networking environments. All VM configuration and statistics remain bound to the VM even when the VM migrates from one physical host in the resource pool to another.

To determine what networking stack is configured, run the following command:

```
1 xe host-list params=software-version
```

In the command output, look for `network_backend`. When the vSwitch is configured as the network stack, the output appears as follows:

```
1 network_backend: openvswitch
```

When the Linux bridge is configured as the network stack, the output appears as follows:

```
1 network_backend: bridge
```

To revert to the Linux network stack, run the following command:

```
1 xe-switch-network-backend bridge
```

Restart your host after running this command.

Citrix Hypervisor networking overview

This section describes the general concepts of networking in the Citrix Hypervisor environment.

Citrix Hypervisor creates a network for each physical NIC during installation. When you add a server 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 Citrix Hypervisor:

- **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 server'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 Citrix Hypervisor servers 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 server 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 Citrix Hypervisor server 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 Citrix Hypervisor server, 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. Citrix Hypervisor servers support VLANs in multiple ways.

Note:

- We recommend not to use a GFS2 SR with a VLAN due to a known issue where you cannot add or remove hosts on a clustered pool if the cluster network is 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 Citrix Hypervisor VLAN features to connect guest virtual network interfaces (VIFs) to specific VLANs. In this case, the Citrix Hypervisor server performs the VLAN tagging/untagging functions for the guest, which is unaware of any VLAN configuration.

Citrix Hypervisor VLANs are represented by additional PIF objects representing VLAN interfaces corresponding to a specified VLAN tag. You can connect Citrix Hypervisor 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 Citrix Hypervisor servers, 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 Citrix Hypervisor server. 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 Citrix Hypervisor 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 and VM networks. 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.

Note:

Citrix Hypervisor supports jumbo frames only when using vSwitch as the network stack on all hosts in the pool.

Requirements for using jumbo frames Note the following when using jumbo frames:

- Jumbo frames are configured at a pool level
- vSwitch must be configured as the network back-end on all servers in the pool
- Every device on the subnet must be configured to use jumbo frames
- Enabling jumbo frames on the management network is not supported

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

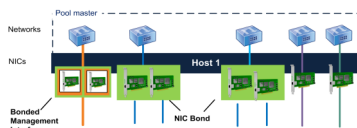
NIC bonds, sometimes also known as NIC teaming, improve Citrix Hypervisor server resiliency and bandwidth by enabling administrators to configure two or more NICs together. NIC bonds logically function as one network card and all bonded NICs share a MAC address.

If one NIC in the bond fails, the host's network traffic is automatically redirected through the second NIC. Citrix Hypervisor supports up to eight bonded networks.

Citrix Hypervisor supports active-active, active-passive, and LACP bonding modes. The number of NICs supported and the bonding mode supported varies according to network stack:

- LACP bonding is only available for the vSwitch whereas active-active and active-passive are available for both the vSwitch and the Linux bridge.
- When the vSwitch is the network stack, you can bond either two, three, or four NICs.
- When the Linux bridge is the network stack, you can only bond two NICs.

In the illustration that follows, the management interface is on a bonded pair of NICs. Citrix Hypervisor uses this bond for management traffic.



All bonding modes support failover. However, not all modes allow all links to be active for all traffic types. Citrix Hypervisor supports bonding the following types of NICs together:

- **NICs (non-management).** You can bond NICs that Citrix Hypervisor 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 a 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 Citrix Hypervisor 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 Citrix Hypervisor, An additional PIF represents a NIC bond. Citrix Hypervisor 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.
- NIC bonds are not supported on NICs that carry FCoE traffic.

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 being provided by different utility companies.
- Consider using uninterruptible power supply units to ensure the network switches and servers can continue to function or can 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. Citrix Hypervisor 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

Citrix Hypervisor 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 server. 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

Citrix Hypervisor 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
```

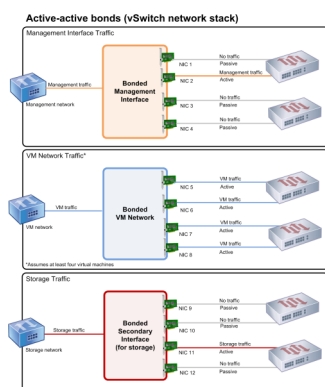
Citrix Hypervisor checks the status of links in bonds approximately every five seconds. Therefore, if more links in the bond fail in the five-second window, the failure is not logged until the next status check.

Bonding event logs appear in the XenCenter Logs tab. 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 when either the Linux bridge or vSwitch network stack is enabled.

When active-active bonding is used with the Linux bridge, you can only bond two NICs. When using the vSwitch as the network stack, 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.



Citrix Hypervisor can only send traffic over two or more NICs when there is more than one MAC address associated with the bond. Citrix Hypervisor 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. Citrix Hypervisor balances the control domain's traffic the same way as it balances VM traffic.

Traffic balancing Citrix Hypervisor 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.

Citrix Hypervisor 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, Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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:

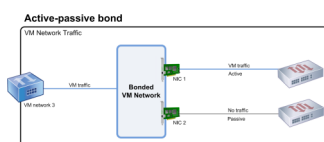
Active-active bonding does not require switch support for EtherChannel or 802.3ad (LACP).

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 is available in the Linux bridge and the vSwitch network stack. When used with the Linux bridge, you can bond two NICs together. When used with the vSwitch, you can only 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 Citrix Hypervisor. 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](#). You must configure the vSwitch as the network stack to use LACP bonding. Also, your switches must support the IEEE 802.3ad standard.

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.
- Supported only for the vSwitch.
- Requires a single switch or stacked switch.

Traffic balancing Citrix Hypervisor 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:

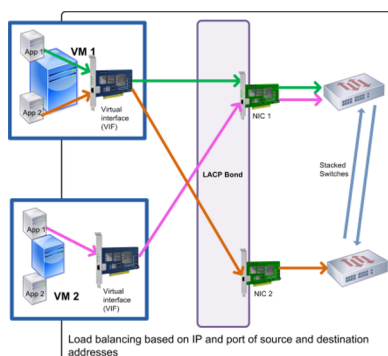
the configuration does not have to match on both sides.

Load balancing based on IP and port of source and destination addresses.

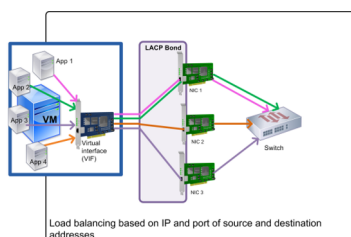
This hashing type is the default LACP bonding hashing algorithm. If there is a variation in the source or destination IP or port numbers, traffic from one guest can be distributed over two links.

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.

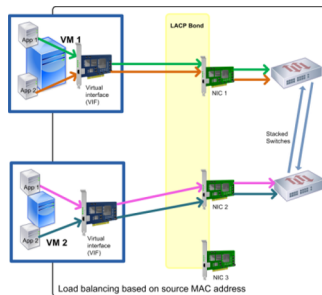


The balancing algorithm for this hashing type uses five factors to spread traffic across the NICs: the source IP address, source port number, destination IP address, destination port number, and source MAC address.

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. Citrix Hypervisor 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:

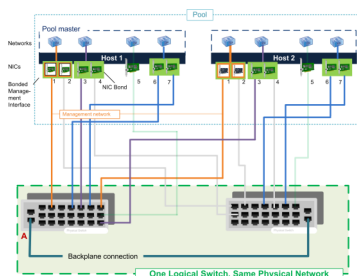
- When you connect one of the links in a bonded management interface to a second switch, if the switch fails, the management network remains online and the hosts can still communicate with each other.
- If you connect a link (for any traffic type) to a second switch and the NIC or switch fails, the virtual 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.

The illustration that follows shows how the cables and network configuration for the bonded NICs have to match.



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 a 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.

- Citrix Hypervisor 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 Citrix Hypervisor server 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 Citrix Hypervisor 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). Additional VLAN network is created (for example, for Pool-wide network associated with eth0 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 Citrix Hypervisor server 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 Citrix Hypervisor 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 master 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 master PIF>
```

```
1 xe pif-plug uuid=<uuid of bond master PIF>
```

Manage networking

May 28, 2024

Network configuration procedures in this section differ depending on whether you are configuring a stand-alone server or a server that is part of a resource pool.

Create networks in a standalone server

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 Citrix Hypervisor server 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 Citrix Hypervisor servers 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 Citrix Hypervisor management traffic.

As all hosts in a pool share a common set of network. It is important to have the same physical networking configuration for Citrix Hypervisor servers in a pool. PIFs on the individual hosts are connected to pool-wide networks based on device name. For example, all Citrix Hypervisor servers in a pool with `eth0` NIC have a corresponding PIF plugged to the pool-wide `Network 0` network. The same is true for hosts with `eth1` NICs and `Network 1`, and other NICs present in at least one Citrix Hypervisor server in the pool.

If one Citrix Hypervisor server 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 eth0 and eth1 are valid on *host2*. VMs on *host1* with VIFs connected to networks corresponding to eth2 and eth3 cannot migrate to host *host2*.

Create VLANs

For servers 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 Citrix Hypervisor server 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 server](#).

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 Citrix Hypervisor servers that are not in a pool. For information on using the xe CLI to create NIC bonds on Citrix Hypervisor servers 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, Citrix Hypervisor 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 Citrix Hypervisor server 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 master and then each member of the pool.
- Using the CLI to configure bonds on the master and then restarting each pool member so that it inherits its settings from the master.
- Using XenCenter to configure the bonds on the master. XenCenter automatically synchronizes the networking settings on the member servers with the master, so you do not need to restart the member servers.

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 Citrix Hypervisor servers 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 master. The master host 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 Citrix Hypervisor 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.

If you are using bonding for resiliency for your storage traffic, you may want to consider using LACP instead of the Linux bridge bonding. To use LACP bonding, you must configure the vSwitch as your networking stack. For more information, see [vSwitch networks](#).

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 Citrix Hypervisor 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 Citrix Hypervisor 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 BIOS 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.
- Only HVM guests are supported with SR-IOV.
- 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 hosts 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.
- 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.

- VMs created in previous releases cannot use this feature from XenCenter.
- 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.
- Hardware restriction: The SR-IOV feature relies on the Controller to reset device functions to a pristine state within 100ms, when requested by the hypervisor using Function Level Reset (FLR).
- 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.

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:kpbs=100
```

Note:

The `kpbs` 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 Citrix Hypervisor server. 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 Citrix Hypervisor server, 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

Citrix Hypervisor servers 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 master and other hosts.

Note:

You must be careful when changing the IP address of a server, 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 server 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 master host by checking that all the other Citrix Hypervisor servers in the pool are visible:

```
1 xe host-list
```

Changing the IP address of the master Citrix Hypervisor server requires extra steps. This is because each pool member uses the advertised IP address of the pool master for communication. The pool members do not know how to contact the master 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 masters.

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 master changes, all member hosts enter into an emergency mode when they fail to contact the master host.

On the pool master, use the `pool-recover-slaves` command to force the master to contact each pool member and inform them of the new master IP address:

```
1 xe pool-recover-slaves
```

Management interface

When you install Citrix Hypervisor on a host, one of its NICs is designated as the *management interface*: the NIC used for Citrix Hypervisor management traffic. The management interface is used for XenCenter and other management API 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 Citrix Hypervisor. For security reasons, you can close TCP port 80 on the management interface. By default, port 80 is still open. If you close it, any external clients that use the management API must use HTTPS over port 443 to connect to Citrix Hypervisor. However, before closing port 80, check whether all your API clients (Citrix Virtual Apps and Desktops in particular) can use HTTPS over port 443.

To 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 Citrix Hypervisor server in the usual manner.
2. Restart your Citrix Hypervisor server.
3. List all the physical NICs for that Citrix Hypervisor server 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 Citrix Hypervisor server 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 Citrix Hypervisor server in the usual manner. After restarting the server, 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 [Citrix Hypervisor Changed Block Tracking Guide](#).

Use switch port locking

The Citrix Hypervisor 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.

Switch-port locking is a feature designed for public cloud-service providers in environments concerned about internal threats. This functionality assists public cloud-service providers who have a network architecture in which each VM has a public, internet-connected IP address. Because cloud tenants are untrusted, you can use security measures such as spoofing protection to ensure that tenants cannot attack other virtual machines in the cloud.

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 Citrix Hypervisor administrator has specified it can use
- Intercepting, spoofing, or disrupting the traffic of other VMs

Requirements

- The Citrix Hypervisor switch-port locking feature is supported on the Linux bridge and vSwitch networking stacks.
- 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 master.

- 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 Citrix Hypervisor 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).

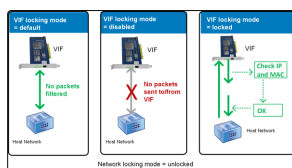
Alice reconfigures Host-B's VIF to be locked to a single MAC but many IP addresses.

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 Citrix Hypervisor 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.

Regardless of which networking stack you use, the feature operates the same way. However, as described in more detail in the sections that follow, the Linux bridge does not fully support switch-port locking in IPv6.

VIF locking-mode states The Citrix Hypervisor 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`, Citrix Hypervisor 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`, Citrix Hypervisor applies a filtering rule so that the VIF drops all traffic.

-`default-locking-mode=unlocked`, Citrix Hypervisor 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.** Citrix Hypervisor 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. However, if you have the Linux bridge configured, do not type IPv6 addresses.

Citrix Hypervisor lets you type IPv6 addresses when the Linux bridge is active. However, Citrix Hypervisor cannot filter based on the IPv6 addresses typed. The reason is the Linux bridge does not have modules to filter Neighbor Discovery Protocol (NDP) packets. Therefore, complete protection cannot be implemented and guests would be able to impersonate another guest by forging NDP packets. As result, if you specify even one IPv6 address, Citrix Hypervisor lets all IPv6 traffic pass through the VIF. If you do not specify any IPv6 addresses, Citrix Hypervisor does not let any IPv6 traffic pass through to the VIF.

- **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, Citrix Hypervisor 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 Citrix Hypervisor 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, Citrix Hypervisor 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 Citrix Hypervisor 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:

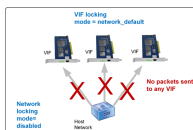
- **Unlocked.** When the network `default-locking-mode` parameter is set to `unlocked`, Citrix Hypervisor lets the VM send traffic to any IP address on the network the VIF connects to.

- **Disabled.** When the **default-locking-mode** parameter is set to **disabled**, Citrix Hypervisor 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 Citrix Hypervisor 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
```

Troubleshoot networking

January 9, 2023

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 apply a BIOS update.

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=eth0
```

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, Citrix Hypervisor server 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 master 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 master requires a network reset, reset the network on the pool master 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 master'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 Master on its new IP address. In this situation, the IP address of the Pool Master 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 Master’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 master, 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 Master’s management interface. Defaults to the last known Pool Master’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 master command-line examples Examples of commands that can be applied on a Pool Master:

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 Master'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 Master's IP address was changed:

```
1 xe-reset-networking --master=master-ip-address
```

To reset networking for Static IP configuration, assuming the Pool Master'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 Master's IP address was changed after initial setup:

```
1 xe-reset-networking --device=device-name --master=master-ip-address
```

Storage

March 16, 2023

This section describes how physical storage hardware maps to virtual machines (VMs), and the software objects used by the management API to perform storage-related tasks. Detailed sections on each of the supported storage types include the following information:

- Procedures for creating storage for VMs using the CLI, with type-specific device configuration options
- Generating snapshots for backup purposes
- Best practices for managing storage

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).

SRs are flexible, with built-in support for the following drives:

Locally connected:

- SATA
- SCSI
- SAS
- NVMe

The local physical storage hardware can be a hard disk drive (HDD) or a solid state drive (SSD).

Remotely connected:

- iSCSI
- NFS
- SAS
- SMB (version 3 only)
- Fibre Channel

Note:

NVMe over Fibre Channel and NVMe over TCP are not supported.

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. This software stack is based on Microsoft's Virtual Hard Disk (VHD) specification.

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 Citrix Hypervisor server 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 servers in a single resource pool must have at least one shared SR in common. 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](#).

Warning:

Citrix Hypervisor does not support snapshots at the external SAN-level of a LUN for any SR type.

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 Citrix Hypervisor. VDIs are persistent, on-disk objects that exist independently of Citrix Hypervisor servers. 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)

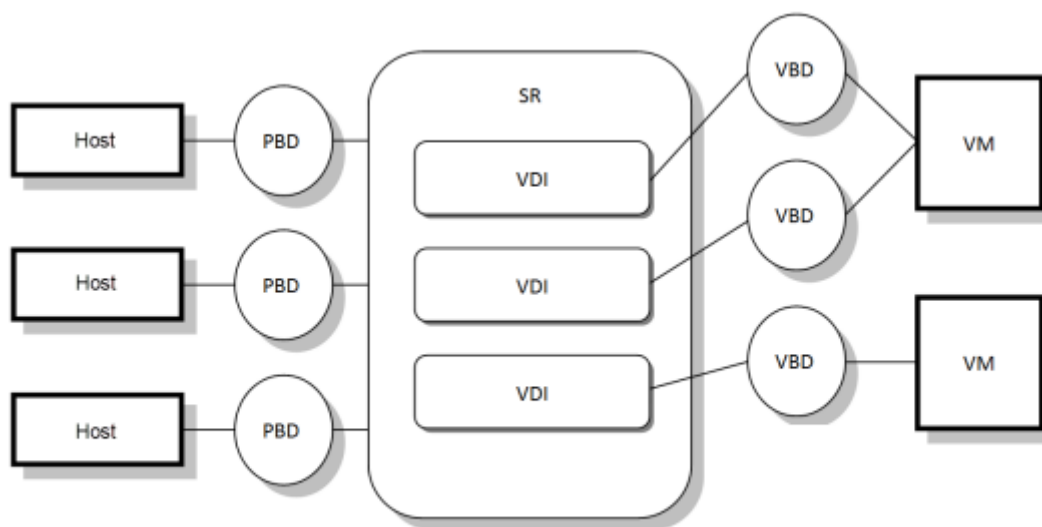
Physical block devices represent the interface between a physical server and an attached SR. PBDs are connector objects that allow a given SR to be mapped to a host. PBDs store the device configuration fields that are used to connect to and interact with a given storage target. For example, NFS device configuration includes the IP address of the NFS server and the associated path that the Citrix Hypervisor server mounts. PBD objects manage the run-time attachment of a given SR to a given Citrix Hypervisor server. 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](#).

Summary of storage objects

The following image is a summary of how the storage objects presented so far are related:



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 Citrix Hypervisor 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).

VDI types

For GFS2 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 a 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 `vdi-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 `vdi-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, Citrix Hypervisor 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 amount of shared data.

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 Read-Write 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. Citrix Hypervisor 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.

VHD-specific notes on coalesce Only one coalescing process is ever active for an SR. This process thread runs on the SR master host.

If you have critical VMs running on the master server of the pool, you can take the following steps to mitigate against occasional slow I/O:

- Migrate the VM to a host other than the SR master

- 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

April 23, 2024

You can use the **New Storage Repository** wizard in XenCenter to create storage repositories (SRs). The wizard guides you through the configuration steps. Alternatively, use the CLI, and the `sr-create` command. The `sr-create` 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 SR `shared = true` flag is set, a PBD record is created and plugged for every Citrix Hypervisor in the resource pool.

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](#).

All Citrix Hypervisor 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. The other SR types (EXT3/EXT4, NFS, GFS2) 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.
- Citrix Hypervisor 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.

The maximum supported VDI sizes are:

Storage Repository Format	Maximum VDI size
EXT3/EXT4	2 TiB
GFS2 (with iSCSI or HBA)	16 TiB
LVM	2 TiB
LVMoFCOE (deprecated)	2 TiB
LVMoHBA	2 TiB
LVMoiSCSI	2 TiB
NFS	2 TiB
SMB	2 TiB

Local LVM

The Local LVM type presents disks within a locally attached Volume Group.

By default, Citrix Hypervisor 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 native blocks, the storage must also support emulation of 512 byte allocation blocks.

LVM performance considerations

The snapshot and fast clone functionality for LVM-based SRs comes with an inherent performance overhead. When optimal performance is required, Citrix Hypervisor supports creation of VDIs in the *raw* format in addition to the default VHD format. The Citrix Hypervisor 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?
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 LVM SR on `/dev/sdb`, 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/sdb 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. 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 Citrix Hypervisor CLI.

Whether a local EXT SR uses EXT3 or EXT4 depends on what version of Citrix Hypervisor created it:

- If you created the local EXT SR on an earlier version of XenServer or Citrix Hypervisor and then upgraded to Citrix Hypervisor 8.2, it uses EXT3.
- If you created the local EXT SR on Citrix Hypervisor 8.2, it uses EXT4.

Note:

The block size of an EXT3/EXT4 disk must be 512 bytes. To use storage with 4 KB native blocks, the storage must also support emulation of 512 byte allocation blocks.

Creating a local EXT4 SR (ext)

Device-config parameters for EXT 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 EXT4 SR on `/dev/sdb`, 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/sdb type=ext
```

udev

The udev type represents devices plugged in using the udev device manager as VDIs.

Citrix Hypervisor 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 Citrix Hypervisor server. The other is for a USB device plugged into a USB port of the Citrix Hypervisor server. VDIs that represent the media come and go as disks or USB sticks are inserted and removed.

ISO

The ISO type handles CD images stored as files in ISO format. This SR type is useful for creating shared ISO libraries.

The following ISO SR types are available:

- `nfs_iso`: The NFS ISO SR type handles CD images stored as files in ISO format available as an NFS share.
- `cifs`: The Windows File Sharing (SMB/CIFS) SR type handles 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, Citrix Hypervisor uses the `location` device config parameter to decide the type.

Device-config parameters for ISO SRs:

Parameter Name	Description	Required?
<code>location</code>	Path to the mount.	Yes

Parameter Name	Description	Required?
<code>type</code>	Storage type to use for the SR: <code>cifs</code> or <code>nfs_iso</code> .	No
<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>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 username 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
<code>cifspassword</code>	For the storage type CIFS/SMB, if a password is required for the Windows file server. We recommend 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 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=<server:/path> device-config:type=nfs_iso
```

You can use NFS or SMB to mount the ISO SR. For more information about using these SR types, see [NFS and SMB](#).

We recommend that you use SMB version 3 to mount ISO SR on Windows file server. Version 3 is selected by default because it is more secure and robust than SMB version 1.0. However, 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=<\\IP\path>  
2  device-config:username=<username> device-config:  
    cifspassword_secret=<password_secret> \  
3  device-config:type=cifs device-config:vers=1.0 name-label="Example  
    ISO SR"
```

Software iSCSI support

Citrix Hypervisor supports shared 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 Citrix Hypervisor server in a resource pool and migrated between them with no noticeable downtime.

iSCSI SRs use the entire LUN specified at creation time and may not 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 native blocks, the storage must also support emulation of 512 byte allocation blocks.

Citrix Hypervisor server iSCSI configuration

All iSCSI initiators and targets must have a unique name to ensure 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.

Citrix Hypervisor servers 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 Citrix Hypervisor server 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 servers in a pool, ensure that multi-initiator access is enabled for the specified LUN.

The Citrix Hypervisor server 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 Citrix Hypervisor server IQN with iSCSI SRs attached. Doing so can result in failures connecting to new targets or existing SRs.

Software FCoE storage (deprecated)

Software FCoE provides a standard framework to which hardware vendors can plug in their FCoE-capable NIC and get the same benefits of a hardware-based FCoE. This feature eliminates the need for using expensive HBAs.

Note:

Software FCoE is deprecated and will be removed in a future release.

Before you create a software FCoE storage, manually complete the configuration required to expose a LUN to the host. This configuration includes configuring the FCoE fabric and allocating LUNs to your SAN's public world wide name (PWWN). After you complete this configuration, the available LUN is mounted to the host's CNA as a SCSI device. The SCSI device can then be used to access the LUN as if it were a locally attached SCSI device. For information about configuring the physical switch and the array to support FCoE, see the documentation provided by the vendor.

Note:

Software FCoE can be used with Open vSwitch and Linux bridge as the network back-end.

Create a Software FCoE SR

Before creating a Software FCoE SR, customers must ensure that there are FCoE-capable NICs attached to the host.

Device-config parameters for FCoE SRs are:

Parameter Name	Description	Required?
<code>SCSIid</code>	The SCSI bus ID of the destination LUN	Yes

Run the following command to create a shared FCoE SR:

```
1  xe sr-create type=lvmofoe \
2  name-label="FCoE SR" shared=true device-config:SCSIid=SCSI_id
```

Hardware host bus adapters (HBAs)

This section covers various operations required to manage SAS, Fibre Channel, and iSCSI HBAs.

Sample QLogic iSCSI HBA setup

For details on configuring QLogic Fibre Channel and iSCSI HBAs, see the [Cavium](#) website.

Once the HBA is physically installed into the Citrix Hypervisor server, use the following steps to configure the HBA:

1. Set the IP networking configuration for the HBA. This example assumes DHCP and HBA port 0. Specify the appropriate values if using static IP addressing or a multi-port HBA.

```
1  /opt/QLogic_Corporation/SANsurferiCLI/isccli -ipdhcp 0
```

2. Add a persistent iSCSI target to port 0 of the HBA.

```
1  /opt/QLogic_Corporation/SANsurferiCLI/isccli -pa 0
    iscsi_target_ip_address
```

3. Use the `xe sr-probe` command to force a rescan of the HBA controller and display available LUNs. For more information, see [Probe an SR](#) and [Create a Shared LVM over Fibre Channel / Fibre Channel over Ethernet / iSCSI HBA or SAS SR](#).

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 Citrix Hypervisor server 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 native blocks, the storage must also support emulation of 512 byte allocation blocks.

Create a shared LVM over iSCSI SR by using the Software iSCSI initiator

Device-config parameters for LVMoiSCSI SRs:

Parameter Name	Description	Required?
<code>target</code>	The IP address or host name of the iSCSI filer that hosts the SR. This can also be a comma-separated list of values.	Yes
<code>targetIQN</code>	The IQN target address of iSCSI filer that hosts the SR	Yes
<code>SCSIid</code>	The SCSI bus ID of the destination LUN	Yes
<code>chapuser</code>	The user name to be used for CHAP authentication	No
<code>chappassword_secret</code>	(Recommended) Secret ID for the password to be used for CHAP authentication. Pass a secret instead of a password.	No
<code>chappassword</code>	The password to be used for CHAP authentication. We recommend you use the <code>chappassword_secret</code> parameter instead.	No
<code>port</code>	The network port number on which to query the target	No
<code>usediscoverynumber</code>	The specific iSCSI record index to use	No
<code>incoming_chapuser</code>	The user name that the iSCSI filter uses to authenticate against the host	No
<code>incoming_chappassword</code>	The password that the iSCSI filter uses to authenticate against the host	No

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 LVMo 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=scsci_id \  
5  type=lvmoiscsi
```

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.

Device-config parameters for LVMoHBA SRs:

Parameter name	Description	Required?
SCSIid	Device SCSI ID	Yes

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 Citrix Hypervisor server in the pool. This process is highly specific to the SAN equipment in use. For more information, see your SAN documentation.
2. If necessary, use the HBA CLI included in the Citrix Hypervisor server to configure the HBA:
 - Emulex: `/bin/sbin/ocmanager`
 - QLogic FC: `/opt/QLogic_Corporation/SANsurferCLI`
 - QLogic iSCSI: `/opt/QLogic_Corporation/SANsurferiCLI`

For an example of QLogic iSCSI HBA configuration, see *Hardware host bus adapters (HBAs)* in the previous section. For more information on Fibre Channel and iSCSI HBAs, see the [Broadcom](#) and [Cavium](#) websites.

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 master host of the pool, 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.

Thin provisioned shared GFS2 block storage

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).

Thin provisioning for shared block storage is of particular interest in the following cases:

- You want increased space efficiency. Images are sparsely and not thickly allocated.
- You want to reduce the number of I/O operations per second on your storage array. The GFS2 SR is the first SR type to support storage read caching on shared block storage.
- You use a common base image for multiple virtual machines. The images of individual VMs will then typically utilize even less space.
- You use snapshots. Each snapshot is an image and each image is now sparse.
- Your storage does not support NFS and only supports block storage. If your storage supports NFS, we recommend you use NFS instead of GFS2.
- You want to create VDIs that are greater than 2 TiB in size. The GFS2 SR supports VDIs up to 16 TiB in size.

Note:

We recommend not to use a GFS2 SR with a VLAN due to a known issue where you cannot add or remove hosts on a clustered pool if the cluster network is on a non-management VLAN.

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.

To use shared GFS2 storage, the Citrix Hypervisor resource pool must be a clustered pool. Enable clustering on your pool before creating a GFS2 SR. For more information, see [Clustered pools](#).

Ensure that storage multipathing is set up between your clustered pool and your GFS2 SR. For more information, see [Storage multipathing](#).

SRs of type GFS2 can be created and managed using the xe CLI or XenCenter.

Create a shared GFS2 SR

You can create your shared GFS2 SR on an iSCSI or an HBA LUN.

Create a shared GFS2 over iSCSI SR

You can create GFS2 over iSCSI SRs by using XenCenter. For more information, see [Software iSCSI storage](#) in the XenCenter product documentation.

Alternatively, 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=ghs2 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 fence.

For more information about working with iSCSI SRs, see [Software iSCSI support](#).

Create a shared GFS2 over HBA SR

You can create GFS2 over HBA SRs by using XenCenter. For more information, see [Hardware HBA storage](#) in the XenCenter product documentation.

Alternatively, you can use the xe CLI to create a GFS2 over HBA SR.

Device-config parameters for GFS2 SRs:

Parameter name	Description	Required?
provider	The block provider implementation. In this case, hba .	Yes
SCSIid	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 host bus adapters](#).

Constraints

Shared GFS2 storage currently has the following constraints:

- 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 or 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 FCoE transport is not supported with GFS2 SRs.
- Trim/unmap is not supported on GFS2 SRs.
- CHAP is not supported on GFS2 SRs.
- MCS full clone VMs are not supported with GFS2 SRs.
- Using multiple GFS2 SRs in the same MCS catalog is not supported.
- Performance metrics are not available for GFS2 SRs and disks on these SRs.
- Changed block tracking is not supported for VDIs stored 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 Citrix Hypervisor to write to it.
- You cannot have more than 62 GFS2 SRs in your pool.
- Clustered pools only support up to 16 hosts per pool.
- To enable HA in a clustered pool, the heartbeat SR must be a GFS2 SR.
- For cluster traffic, you must 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 clustering network while the cluster is live and has running VMs. This action can cause the cluster to fence.
- If you have an IP address conflict (multiple hosts having the same IP address) on your clustering network involving at least one host with clustering enabled, the hosts do not fence. To fix this issue, resolve the IP address conflict.

NFS and SMB

Shares on NFS servers (that support any version of NFSv4 or NFSv3) or 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. Additionally, as these SRs can be shared, VDIs stored on shared SRs allow:

- VMs to be started on any Citrix Hypervisor servers in a resource pool
- VM migrate between Citrix Hypervisor servers 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 and are not supported in Citrix Hypervisor 8.2.
- Clustered SMB is not supported with Citrix Hypervisor.
- For NFSv4, only the authentication type `AUTH_SYS` is supported.
- SMB storage is available for Citrix Hypervisor Premium Edition customers, or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.
- It is highly recommended for both NFS and SMB storage that a dedicated storage network be used, 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 very fast provisioning and deployment of new VMs.

Note:

The maximum supported length of VHD chains is 30.

File-based SRs and VHD implementations in Citrix Hypervisor 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.

Citrix Hypervisor 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 fail-

ure. Citrix Hypervisor has been tested extensively against Network Appliance FAS2020 and FAS3210 storage, using Data OnTap 7.3 and 8.1

Warning:

As VDIs on file-based SRs are created as thin provisioned, administrators must ensure that the file-based SRs have enough disk space for all required VDIs. Citrix Hypervisor servers 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.

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 hostname 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 Citrix Hypervisor 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 servers in the pool. If this configuration is not done, the creation of the SR and the plugging of the PBD record fails.

The Citrix Hypervisor NFS implementation uses TCP by default. If your situation allows, you can configure the implementation to use UDP in scenarios where there may be a performance benefit. To do this configuration, when creating an SR, specify the `device-config` parameter `useUDP=true`.

The following `device-config` parameters are used with NFS SRs:

Parameter Name	Description	Required?
<code>server</code>	IP address or hostname of the NFS server	Yes

Parameter Name	Description	Required?
<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

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"
```

Create a shared SMB SR (SMB)

To create an SMB SR, provide the hostname or IP address of the SMB server, the full path of the exported share, and appropriate credentials.

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

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
```

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.

Citrix Hypervisor servers 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 Citrix Hypervisor server. 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.

See [Create storage repositories](#) for details on creating shared HBA-based FC and iSCSI SRs.

Note:

Citrix Hypervisor 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 native blocks, the storage must also support emulation of 512 byte allocation blocks.

Thin-provisioned shared GFS2 block storage

April 29, 2024

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).

Thin provisioning for shared block storage is of particular interest in the following cases:

- You want increased space efficiency. Images are sparsely and not thickly allocated.
- You want to reduce the number of I/O operations per second on your storage array. The GFS2 SR is the first SR type to support storage read caching on shared block storage.
- You use a common base image for multiple virtual machines. The images of individual VMs will then typically utilize even less space.
- You use snapshots. Each snapshot is an image and each image is now sparse.
- You want to create VDIs that are greater than 2 TiB in size. The GFS2 SR supports VDIs up to 16 TiB in size.

- Your storage doesn't support NFS or SMB3 and only supports block storage. If your storage supports NFS or SMB3, we recommend you use these SR types instead of GFS2.
- Your storage doesn't support thin provisioning of LUNs. If your storage does thin provision LUNs, you can encounter problems and run out of space when combining it with GFS2. Combining GFS2 with a thin-provisioned LUN does not provide many additional benefits and is not recommended.

Note:

We recommend not to use a GFS2 SR with a VLAN due to a known issue where you cannot add or remove hosts on a clustered pool if the cluster network is on a non-management VLAN.

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, Citrix Hypervisor requires that you use a clustered pool with your GFS2 SR. We also recommend that you design your environment and configure Citrix Hypervisor features to provide as much resiliency and redundancy as possible.

Before setting up your Citrix Hypervisor 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.

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, 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 clustering 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 Citrix Hypervisor](#).

2. Open a console on the Citrix Hypervisor server that you want to act as the pool master.
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 master, when you join other Citrix Hypervisor servers 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 clustering network while the cluster is live and has run-

- ning 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 clustering 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.

To test your active-passive bonded network failover times:

For bonded networks that use active-passive mode, if the active link fails, there is a failover period when the network link is broken while the passive link becomes active. If the time it takes for your active-passive bonded network to fail over is longer than the cluster timeout, some or all hosts in your clustered pool might still fence.

You can test your bonded network failover time by forcing the network to fail over by using one of the following methods:

- By physically pulling out the network cables
- By disabling switch ports on one network link

Repeat the test a number of times to ensure the result is consistent.

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 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 Citrix Hypervisor resource pool must be a clustered pool. Enable clustering on your pool before creating a GFS2 SR.

A clustered pool is a pool of Citrix Hypervisor servers 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 Citrix Hypervisor servers 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 Citrix Hypervisor servers.

Repeat the following steps on each joining Citrix Hypervisor server that is not the pool master:

- a) Open a console on the Citrix Hypervisor server.
- b) Join the Citrix Hypervisor server to the pool on the pool master 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 Citrix Hypervisor server that is the pool master. The `password` must be the administrator password set when the pool master 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 Citrix Hypervisor server 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 Citrix Hypervisor server 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 Citrix Hypervisor server.
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 Citrix Hypervisor servers 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 Citrix Hypervisor 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 Citrix Hypervisor. 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 create GFS2 over iSCSI SRs by using XenCenter. For more information, see [Software iSCSI storage](#) in the XenCenter product documentation.

Alternatively, 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 support](#).

Create a shared GFS2 over HBA SR

You can create GFS2 over HBA SRs by using XenCenter. For more information, see [Hardware HBA storage](#) in the XenCenter product documentation.

Alternatively, 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 host bus adapters](#).

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:

- 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.
- MCS full clone VMs are not supported with GFS2 SRs.
- Using multiple GFS2 SRs in the same MCS catalog is not supported.
- Performance metrics are not available for GFS2 SRs and disks on these SRs.
- Changed block tracking is not supported for VDIs stored 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 Citrix Hypervisor 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 Citrix Hypervisor 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 clustering 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 clustering 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

March 4, 2024

This section covers creating storage repository types and making them available to your Citrix Hypervisor server. It also covers various operations required in the ongoing management of Storage Repositories (SRs), including Live VDI Migration.

Create storage repositories

This section explains how to create Storage Repositories (SRs) of different types and make them available to your Citrix Hypervisor server. The examples provided cover creating SRs using the xe CLI. For details on using the **New Storage Repository** wizard to add SRs using XenCenter, see the [XenCenter documentation](#).

Note:

Local SRs of type `lvm` and `ext` can only be created using the xe CLI. After creation, you can manage all SR types by either XenCenter or the xe CLI.

There are two basic steps to create a storage repository for use on a host by using the CLI:

1. Probe the SR type to determine values for any required parameters.
2. Create the SR to initialize the SR object and associated PBD objects, plug the PBDs, and activate the SR.

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.

SRs can be *destroyed* when no longer in use to free up the physical device. SRs can also be *forgotten* to detach the SR from one Citrix Hypervisor server and attach it to another. For more information, see *Removing SRs* in the following section.

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
5      target IQN parameter, \
6      <?xml version="1.0" ?>
7      <iscsi-target-iqns>
8          <TGT>
9              <Index>
10                 0
11             </Index>
12             <IPAddress>
13                 192.168.1.10
14             </IPAddress>
15             <TargetIQN>
16                 iqn.192.168.1.10:filer1
17             </TargetIQN>
18         </TGT>
19     </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             1494554000000000000000000000000000000000000000000000000b70200000f000000
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=149455400000000000000000002000000
   b70200000f000000
4
5  <?xml version="1.0" ?>
6  <SRlist>
7    <SR>
8      <UUID>
9        3f6e1ebd-8687-0315-f9d3-b02ab3adc4a6
10     </UUID>
11     <Devlist>
12       /dev/disk/by-id/scsi-149455400000000000000000002000000
        b70200000f000000
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

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 Citrix Hypervisor server, 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 Citrix Hypervisor server:

```
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 Citrix Hypervisor server 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 Citrix Hypervisor server 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 Citrix Hypervisor servers 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=14945540000000000000000000000020000000  
   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=14945540000000000000000000000020000000  
   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 server 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 Citrix Hypervisor server. 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.

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 Citrix Hypervisor 8.2.
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 server 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 **Reclaim 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 Citrix Hypervisor, 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. However, GFS2 SRs cannot perform leaf coalesce - that is, the VDI that the VM is writing to cannot be coalesced with its parent on a GFS2 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/re-sume 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 master (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 Citrix Hypervisor 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 master.

If the VDIs to be coalesced are on local storage, run the off-line coalesce tool on the server 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 EqualLogic, NetApp, or NFS storage.

Adjust the disk I/O scheduler

For general performance, the default disk scheduler `noop` is applied on all new SR types. The `noop` 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: `noop`, `cfq`, or `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 `cfq` 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 `cfq` 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=cfq
5 xe pbd-unplug uuid=<pbd_uuid>
6 xe pbd-plug uuid=<pbd_uuid>
```

Storage multipathing

April 29, 2024

Dynamic multipathing support is available for Fibre Channel and iSCSI storage back-ends.

Citrix Hypervisor uses Linux native multipathing (DM-MP), the generic Linux multipathing solution, as its multipath handler. However, Citrix Hypervisor supplements this handler with additional features so that Citrix Hypervisor 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 servers in your pool *before* creating the SR. If you create the SR before enabling multipathing, you must put your servers 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 multipathed 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 server and choose **Enter Maintenance Mode**.
2. Wait until the server reappears in the **Resources** pane with the maintenance mode icon (a blue square) before continuing.
3. On the **General** tab for the server, 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 server and choose **Exit Maintenance Mode**.
7. Repeat these steps to enable multipathing on all servers in the pool.

Ensure that you enable multipathing on all servers in the pool. All cabling and, in the case of iSCSI, subnet configurations must match the corresponding NICs on each server.

To enable multipathing by using the xe CLI

1. Open a console on the Citrix Hypervisor server.

2. Unplug all PBDs on the server 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 servers in the pool.

Ensure that you enable multipathing on all servers in the pool. All cabling and, in the case of iSCSI, subnet configurations must match the corresponding NICs on each server.

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 server and choose **Enter Maintenance Mode**.
2. Wait until the server reappears in the **Resources** pane with the maintenance mode icon (a blue square) before continuing.
3. On the **General** tab for the server, 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 server and choose **Exit Maintenance Mode**.
7. Repeat these steps to configure multipathing on all servers in the pool.

To disable multipathing by using the xe CLI

1. Open a console on the Citrix Hypervisor server.
2. Unplug all PBDs on the server 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 servers in the pool.

Configure multipathing

To do additional multipath configuration, 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 Citrix Hypervisor.

Multipath tools

Multipath support in Citrix Hypervisor is based on the device-mapper `multipathd` components. The Storage Manager API handles activating and deactivating multipath nodes automatically. Unlike the standard `dm-multipath` tools in Linux, device mapper nodes are not automatically created for all LUNs on the system. Device mapper nodes are only provisioned when LUNs are actively used by the storage management layer. Therefore, it is unnecessary to use any of the `dm-multipath` CLI tools to query or refresh DM table nodes in Citrix Hypervisor.

If it is necessary to query the status of device-mapper tables manually, or list active device mapper multipath nodes on the system, use the `mpathutil` utility:

```
1 mpathutil list
```

```
1 mpathutil status
```

iSCSI multipathing on a single subnet

You can configure Citrix Hypervisor 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, Citrix Hypervisor 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 Citrix Hypervisor.

To configure multipath:

1. Back up any data you want to protect.
2. In the XenCenter **Resources** pane, right-click on the server and choose **Enter Maintenance Mode**.
3. Wait until the server reappears in the **Resources** pane with the maintenance mode icon (a blue square) before continuing.
4. On the **General** tab for the server, 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 server 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 server and choose **Exit Maintenance Mode**. Do not resume your VMs yet.

9. In the server 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 servers 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 server 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.

IntelliCache

December 20, 2023

Note:

This feature is only supported when using Citrix Hypervisor with Citrix Virtual Desktops.

Intellicache is not supported for VMs using a GFS2 SR.

Using Citrix Hypervisor 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 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 master image from shared storage.

IntelliCache works by caching data from a VMs parent VDI in local storage on the VM host. This local cache is then populated as data is read from the parent VDI. When many VMs share a common parent VDI, a VM can use the data read into the cache from another VM. Further access to the master image on shared storage is not required.

A thin provisioned, local SR is required for IntelliCache. 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.

Important:

Thin Provisioning changes the default local storage type of the host from LVM to EXT4. Thin Pro-

visioning **must be** enabled in order for Citrix Virtual Desktops local caching to work properly.

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.

Warning:

Thin provisioned SRs may 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

IntelliCache must be enabled either during host installation or be enabled manually on a running host using the CLI.

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. The shared storage type, used to host the source Virtual Disk Image (VDI), must be NFS or EXT3/EXT4 based.

Enable on host installation

To enable IntelliCache 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
```

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 located 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 “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 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: A VDI cache file is only deleted when the VDI itself is deleted. The cache is reset when a VDI is attached to a VM (for example on VM start). If the host is offline when you delete the VDI, the SR synchronization that runs on startup garbage collects the cache file.

Note:

The cache file is not deleted from the host when a VM migrates to a different host or is shut down.

Storage read caching

July 15, 2024

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 Citrix Hypervisor Premium Edition customers.

Storage Read Caching is also available for customers who access Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.

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.

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 [CTX220763](#).

- 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 Citrix Hypervisor'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 Citrix Hypervisor server, 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 Citrix Hypervisor Control Domain (dom0) is 64-bit, large values can be used, for example 32768 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 Citrix Hypervisor 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.

PVS-Accelerator

May 28, 2024

The Citrix Hypervisor PVS-Accelerator feature offers extended capabilities for customers using Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor servers. 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 Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. To use the PVS-Accelerator feature, upgrade the License Server to version 11.14.

To use PVS-Accelerator with UEFI-enabled VMs, ensure that you are using Citrix Provisioning 1906 or later.

After upgrading the PVS-Accelerator supplemental pack, XenCenter might list multiple versions of the PVS-Accelerator. However, only the latest version is active. There is no need to uninstall PVS-Accelerator, as old versions of this feature are always superseded by the newest version.

How does PVS-Accelerator work

PVS-Accelerator employs a Proxy mechanism that resides in the Control Domain (dom0) of Citrix Hypervisor. When this feature is enabled, Citrix Provisioning targets device (VM) read requests are cached directly on the Citrix Hypervisor server machine. These requests are cached in physical memory or a storage repository. When subsequent VMs on that Citrix Hypervisor server 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:

- The PVS-Accelerator user interfaces in XenCenter and Citrix Provisioning are only exposed if the PVS-Accelerator supplemental pack is installed.
- Citrix Provisioning target devices are aware of their proxy status. No additional configuration is required once the capability is installed.
- 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 Citrix Hypervisor server.
- 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 Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Desktops and Citrix Virtual Apps entitlement or Citrix DaaS entitlement.
- PVS-Accelerator requires License Server 11.14 or later.
- PVS-Accelerator uses capabilities of OVS and is therefore not available on hosts that use Linux Bridge as the network back-end.
- 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 Citrix Hypervisor and in Citrix Provisioning to enable the PVS-Accelerator feature:

1. Install the PVS-Accelerator Supplemental Pack on each Citrix Hypervisor server in the pool. The supplemental pack is available to download from the [Citrix Hypervisor Product Downloads](#) page. You can install the supplemental pack using XenCenter or the xe CLI. For information about installing a supplemental pack using XenCenter, see [Installing Supplemental Packs](#) in the XenCenter documentation. For CLI instructions, see the [Citrix Hypervisor Supplemental Packs and the DDK Guide](#).
2. Configure PVS-Accelerator in Citrix Hypervisor 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 Citrix Hypervisor by using the CLI* in the following section.
 - For information about configuring PVS-Accelerator using XenCenter, see [PVS-Accelerator](#) in the XenCenter documentation.
3. After configuring PVS-Accelerator in Citrix Hypervisor, 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 Citrix technologies](#).

The configured port range in Citrix Hypervisor 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 Citrix Hypervisor by using the CLI

1. Run the following command to create a Citrix Provisioning site configuration on Citrix Hypervisor:

```
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 Citrix Hypervisor, 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 Citrix Hypervisor 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 **Citrix Hypervisor** screen appears on the Streamed Virtual Machine Setup wizard. It displays the list of all Citrix Provisioning sites configured on Citrix Hypervisor 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 Citrix Hypervisor server.
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 Citrix Hypervisor. 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 Citrix Hypervisor 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
```

Graphics overview

October 26, 2023

This section provides an overview of the virtual delivery of 3D professional graphics applications and workstations in Citrix Hypervisor. The offerings include GPU Pass-through (for NVIDIA, AMD and Intel GPUs) and hardware-based GPU sharing with NVIDIA vGPU™, AMD MxGPU™, and Intel GVT-g™.

Graphics Virtualization is available for Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. To learn more about Citrix Hypervisor editions, and to find out how to upgrade, visit the [Citrix 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.

Citrix Hypervisor allows you to assign a physical GPU (in the Citrix Hypervisor server) to a Windows or HVM Linux VM running on the same host. This GPU pass-through feature is intended for graphics power users, such as CAD designers.

Shared GPU

Shared GPU 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 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.

Vendor support

The following table lists guest support for the GPU, shared GPU, and multiple vGPU features:

	GPU pass-through for Windows VMs	GPU pass-through for HVM 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		YES			
Intel	YES		YES			
NVIDIA	YES	YES	YES	YES	YES (see note)	YES (see note)

Note:

- 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 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 of VMs with vGPU enabled from previous versions of Citrix Hypervisor or XenServer to Citrix Hypervisor 8.2 is not supported.
- 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.

Vendor support

The following table lists support for vGPU live migration:

				Multiple shared GPU (vGPU) for Windows VMs	Multiple shared GPU (vGPU) for Linux VMs
GPU pass- through for Windows VMs	GPU pass- through for HVM Linux VMs	Shared GPU (vGPU) for Windows VMs	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

Citrix Hypervisor 8.2 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 Server 2016 (64-bit) *
- Windows Server 2019 (64-bit) *
- Windows Server 2022 (64-bit) *

Linux guests:

- RHEL 7 *
- RHEL 8 *
- RHEL 9 *
- CentOS 7
- CentOS Stream 9
- Ubuntu 20.04 *
- Rocky Linux 8 *
- Rocky Linux 9 *

AMD MxGPU

Windows guests:

- Windows 10 (64-bit)
- Windows Server 2016 (64-bit)
- Windows Server 2019 (64-bit)

Intel GVT-g

Windows guests:

- Windows 10 (64-bit)
- Windows Server 2016 (64-bit)

Constraints

- VMs with a virtual GPU are not supported with Dynamic Memory Control.

- Citrix Hypervisor 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, Intel GVT-d, Intel GVT-G, AMD MxGPU, 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

July 22, 2024

This section provides step-by-step instructions on how to prepare Citrix Hypervisor for supported graphical virtualization technologies. The offerings include NVIDIA vGPU, AMD MxGPU, and Intel GVT-d and GVT-g.

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 Citrix Hypervisor 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.

To ensure that you always have the latest security and functional fixes, ensure that you install any updates provided by NVIDIA for the drivers on your VMs and the NVIDIA Virtual GPU Manager running in your host server.

Important:

If you are using NVIDIA A16/A2 cards, ensure that you have the following files installed on your Citrix Hypervisor 8.2 hosts:

- The latest version of the NVIDIA host driver (NVIDIA-vGPU-CitrixHypervisor-8.2-535.42.x86_64 or later)
- [Hotfix XS82ECU1027](#) (or a later hotfix that includes XS82ECU1027)

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 Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. To learn more about Citrix Hypervisor editions, and to find out how to upgrade, visit the [Citrix website](#). For more information, see [Licensing](#).

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.

For example, a TESLA M60 card has two physical GPUs, and can support 11 types of vGPU:

- GRID M60-1A
- GRID M60-2A
- GRID M60-4A
- GRID M60-8A
- GRID M60-0B
- GRID M60-1B
- GRID M60-0Q
- GRID M60-1Q
- GRID M60-2Q

- GRID M60-4Q
- GRID M60-8Q

In the case where you start both a VM that has vGPU type M60-1A and a VM that has vGPU type M60-2A:

- One physical GPU only supports M60-1A instances
- The other only supports M60-2A instances

You cannot start any M60-4A instances on that single card.

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](#).
- Citrix Hypervisor Premium Edition (or access to Citrix Hypervisor through a Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement).
- A server capable of hosting Citrix Hypervisor and the supported NVIDIA cards.

Note:

Some NVIDIA GPUs do not support hosts with more than 1 TB of memory. If you are using the following GPUs based on the Maxwell architecture: Tesla M6, Tesla M10, and Tesla M60, ensure that your server has less than 1 TB memory. For more information, see the [NVIDIA documentation](#).

In general, we recommend that for NVIDIA vGPUs, you use a server with less than 768 GB of memory.

- NVIDIA vGPU software package for Citrix Hypervisor, consisting of the NVIDIA Virtual GPU Manager for Citrix Hypervisor, and NVIDIA drivers.
- To run Citrix Virtual Desktops with VMs running NVIDIA vGPU, you also need: Citrix Virtual Desktops 7.6 or later, full installation.

Note:

Review the NVIDIA Virtual GPU User Guide (Ref: DU-06920-001) available from the [NVIDIA website](#). Register with NVIDIA to access these components.

- For NVIDIA Ampere vGPUs and all future generations, you have to enable SR-IOV in your system BIOS.

vGPU live migration

Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor
2. Install the NVIDIA Virtual GPU Manager for Citrix Hypervisor
3. Restart the Citrix Hypervisor server

Installation on Citrix Hypervisor

Citrix Hypervisor is available for download from the [Citrix Hypervisor Downloads](#) page.

Install the following:

- **Citrix Hypervisor Base Installation ISO**
- **XenCenter Windows Management Console**

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

Licensing note

vGPU is available for Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. To learn more about Citrix Hypervisor editions, and to find out how to upgrade, visit the [Citrix 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 Citrix Hypervisor

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 Citrix Hypervisor Control Domain (dom0). It is provided as either a supplemental pack or an RPM file. For more information about installation, see the User Guide included in the NVIDIA vGPU Software.

The Update can be installed in one of the following methods:

- Use XenCenter (**Tools > Install Update > Select update or supplemental pack from disk**)
- Use the xe CLI command `xe-install-supplemental-pack`.

Note:

If you are installing the NVIDIA Virtual GPU Manager using an RPM file, ensure that you copy the RPM file to dom0 and then install.

1. Use the rpm command to install the package:

```
1 rpm -iv <vgpu_manager_rpm_filename>
```

2. Restart the Citrix Hypervisor server:

```
1 shutdown -r now
```

3. After you restart the Citrix Hypervisor server, verify that the software has been installed and loaded correctly by checking the NVIDIA kernel driver:

```
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.
8 Fan Temp Perf Pwr:Usage/Cap| Memory-Usage | GPU-Util
9 Compute M. |
10 +-----+
11 | 0 Tesla M60 On | 0000:05:00.0 Off | Off |
12 | N/A 33C P8 24W / 150W | 7249MiB / 8191MiB | 0%
13 | Default |
14 +-----+
15 | 1 Tesla M60 On | 0000:09:00.0 Off | Off |
16 | N/A 36C P8 24W / 150W | 7249MiB / 8191MiB | 0%
17 | Default |
18 +-----+
19 | 2 Tesla M60 On | 0000:85:00.0 Off | Off |
20 | N/A 36C P8 23W / 150W | 19MiB / 8191MiB | 0%
21 | Default |
22 +-----+
23 | 3 Tesla M60 On | 0000:89:00.0 Off | Off |
24 | N/A 37C P8 23W / 150W | 14MiB / 8191MiB | 0%
25 | Default |
26 +-----+
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
```

AMD MxGPU

AMDs MxGPU enables multiple Virtual Machines (VM) to have direct access to a portion of a single physical GPU, using Single Root I/O Virtualization. The same AMD graphics driver deployed on non-

virtualized operating systems can be used inside the guest.

VMs use MxGPU GPUs in the same manner as a physical GPU that the hypervisor has passed through. An AMD graphics driver loaded in the VM provides direct access to the GPU for performance critical fast paths.

To ensure that you always have the latest security and functional fixes, ensure that you install any updates provided by AMD for the drivers on your VMs.

For more information about using AMD MxGPU with Citrix Hypervisor, see the [AMD Documentation](#).

Licensing note

MxGPU is available for Citrix Hypervisor Premium Edition customers, or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. To learn more about Citrix Hypervisor editions, and to find out how to upgrade, visit the [Citrix website](#). For detailed information on Licensing, see the [Citrix Hypervisor Licensing FAQ](#).

Available AMD MxGPU vGPU types

AMD MxGPU cards can contain multiple GPUs. For example, S7150 cards contain one physical GPU and S7150x2 cards contain two GPUs. Each physical GPU can host several different types of virtual GPU (vGPU). vGPU types split a physical GPU into a pre-defined number of vGPUs. Each of these vGPUs has an equal share of the frame buffer and graphics processing abilities. The different vGPU types are targeted at different classes of workload. vGPU types that split a physical GPU into fewer pieces are more suitable for intensive workloads.

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 on physical GPUs on the same card. This restriction is automatic and can cause unexpected capacity planning issues.

AMD MxGPU system requirements

- AMD FirePro S7100-series GPUs
- Citrix Hypervisor Premium Edition (or access to Citrix Hypervisor through a Citrix Virtual Desktops or Citrix Virtual Apps entitlement or Citrix DaaS entitlement)
- A server capable of hosting Citrix Hypervisor and AMD MxGPU cards. The list of servers validated by AMD can be found on [the AMD website](#).

- AMD MxGPU host drivers for Citrix Hypervisor. These drivers are available from [the AMD download site](#).
- AMD FirePro in-guest drivers, suitable for MxGPU on Citrix Hypervisor. These drivers are available from [the AMD download site](#).
- To run Citrix Virtual Desktops with VMs running AMD MxGPU, you also need Citrix Virtual Desktops 7.13 or later, full installation.
- System BIOS configured to support SR-IOV and the MxGPU configured as the secondary adapter

Preparation overview

1. Install Citrix Hypervisor
2. Install the AMD MxGPU host drivers for Citrix Hypervisor
3. Restart the Citrix Hypervisor server

Installation on Citrix Hypervisor

Citrix Hypervisor is available for download from the [Citrix Hypervisor Downloads](#) page.

Install the following:

- **Citrix Hypervisor 8.2 Cumulative Update**
- **XenCenter 8.2 Windows Management Console**

For more information about installation, see the [Citrix Hypervisor Installation Guide](#).

Install the AMD MxGPU host driver for Citrix Hypervisor

Complete the following steps to install the host driver.

1. The update that contains the driver can be installed by using XenCenter or by using the xe CLI.
 - To install by using XenCenter, go to **Tools > Install Update > Select update or supplemental pack from disk**
 - To install by using the xe CLI, copy the update to the host and run the following command in the directory where the update is located:

```
1 xe-install-supplemental-pack mxgpu-1.0.5.amd.iso
```

2. Restart the Citrix Hypervisor server.

3. After restarting the Citrix Hypervisor server, verify that the MxGPU package has been installed and loaded correctly. Check whether the `gim` kernel driver is loaded by running the following commands in the Citrix Hypervisor server console:

```
1 modinfo gim
2 modprobe gim
```

4. Verify that the `gim` kernel driver has successfully created MxGPU Virtual Functions, which are provided to the guests. Run the following command:

```
1 lspci | grep "FirePro S7150"
```

The output from the command shows Virtual Functions that have the “S7150V” identifier.

5. Use the GPU tab in XenCenter to confirm that MxGPU Virtual GPU types are listed as available on the system.

After the AMD MxGPU drivers are installed, the **Passthrough** option is no longer available for the GPUs. Instead use the **MxGPU.1** option for pass-through.

The following options are also supported: **MxGPU.2** and **MxGPU.4**.

Create a MxGPU enabled VM

Before configuring a VM to use MxGPU, install the VM. Ensure that AMD MxGPU supports the VM operating system. For more information, see [Guest support and constraints](#).

After the VM is installed, complete the configuration by following the instructions in [Create vGPU enabled VMs](#).

Intel GVT-d and GVT-g

Citrix Hypervisor supports Intel’s virtual GPU (GVT-g), a graphics acceleration solution that requires no additional hardware. It uses the Intel Iris Pro feature embedded in certain Intel processors, and a standard Intel GPU driver installed within the VM.

To ensure that you always have the latest security and functional fixes, ensure that you install any updates provided by Intel for the drivers on your VMs and the firmware on your host server.

Intel GVT-d and GVT-g are compatible with the HDX 3D Pro features of Citrix Virtual Apps and Desktops or Citrix DaaS. For more information, see [HDX 3D Pro](#).

Note:

Because the Intel Iris Pro graphics feature is embedded within the processors, CPU-intensive applications can cause power to be diverted from the GPU. As a result, you might not experience

full graphics acceleration as you do for purely GPU-intensive workloads.

Intel GVT-g system requirements and configuration

To use Intel GVT-g, your Citrix Hypervisor server must have the following hardware:

- A CPU that has Iris Pro graphics. This CPU must be listed as supported for Graphics on the [Hardware Compatibility List](#)
- A motherboard that has a graphics-enabled chipset. For example, C226 for Xeon E3 v4 CPUs or C236 for Xeon E3 v5 CPUs.

Note:

Ensure that you restart the hosts when switching between Intel GPU pass-through (GVT-d) and Intel Virtual GPU (GVT-g).

When configuring Intel GVT-g, the number of Intel virtual GPUs supported on a specific Citrix Hypervisor server depends on its GPU bar size. The GPU bar size is called the ‘Aperture size’ in the BIOS. We recommend that you set the Aperture size to 1,024 MB to support a maximum of seven virtual GPUs per host.

If you configure the Aperture size to 256 MB, only one VM can start on the host. Setting it to 512 MB can result in only three VMs being started on the Citrix Hypervisor server. An Aperture size higher than 1,024 MB is not supported and **does not** increase the number of VMs that start on a host.

Enable Intel GPU pass-through

Citrix Hypervisor supports the GPU pass-through feature for Windows VMs using an Intel integrated GPU device.

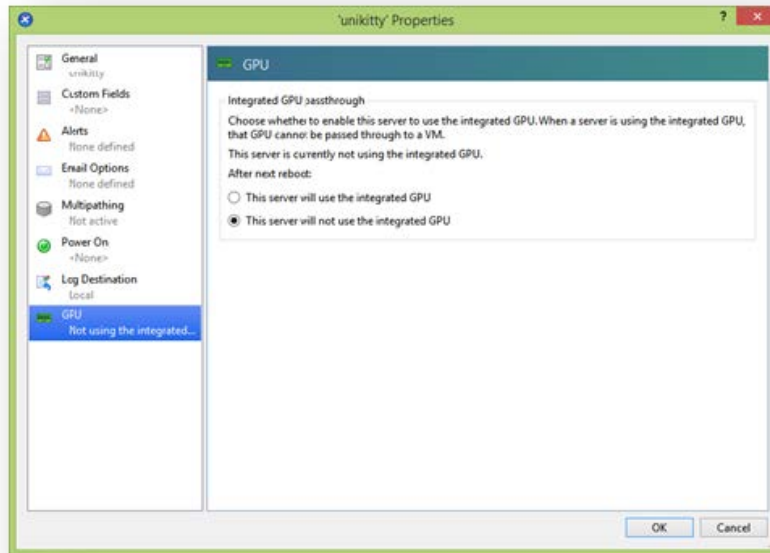
- For more information on Windows versions supported with Intel GPU pass-through, see [Graphics](#).
- For more information on supported hardware, see the [Hardware Compatibility List](#).

When using Intel GPU on Intel servers, the Citrix Hypervisor 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 Citrix Hypervisor server.
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 Citrix Hypervisor server 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 Citrix Hypervisor server'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

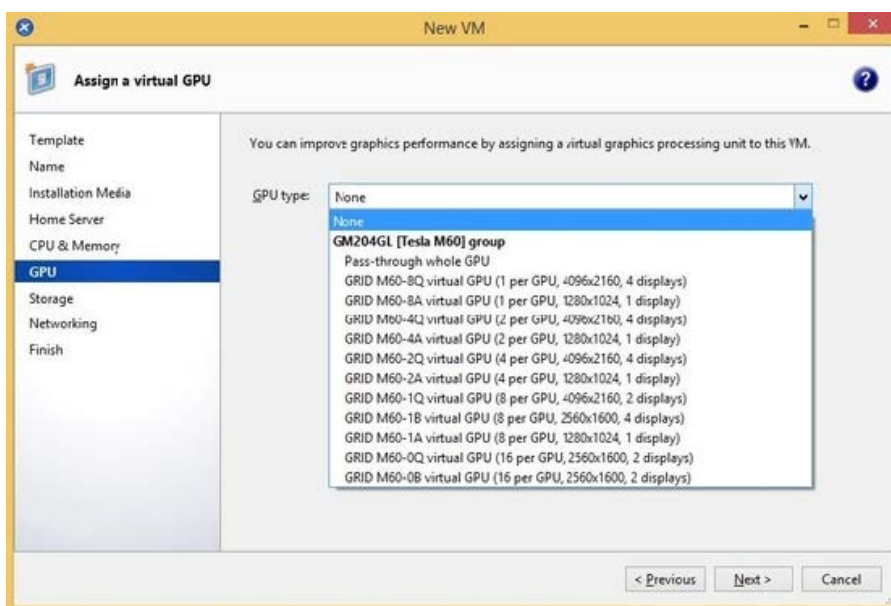
November 27, 2023

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 servers 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 (formerly Citrix 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 Citrix VM Tools for Linux on your VM. For more information, see [Install Citrix 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 or MxGPU, 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.

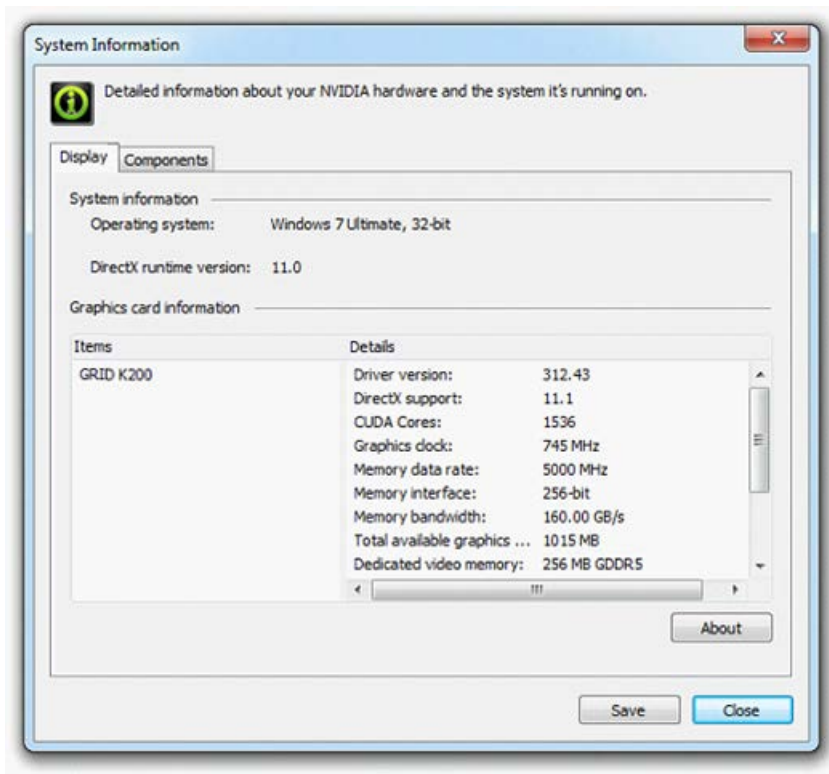
Install the NVIDIA drivers

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 User Guides.

1. Start the VM. In the **Resources** pane, right-click on the VM, and click **Start**.
During this start process, Citrix Hypervisor 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 32-bit or 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](#).

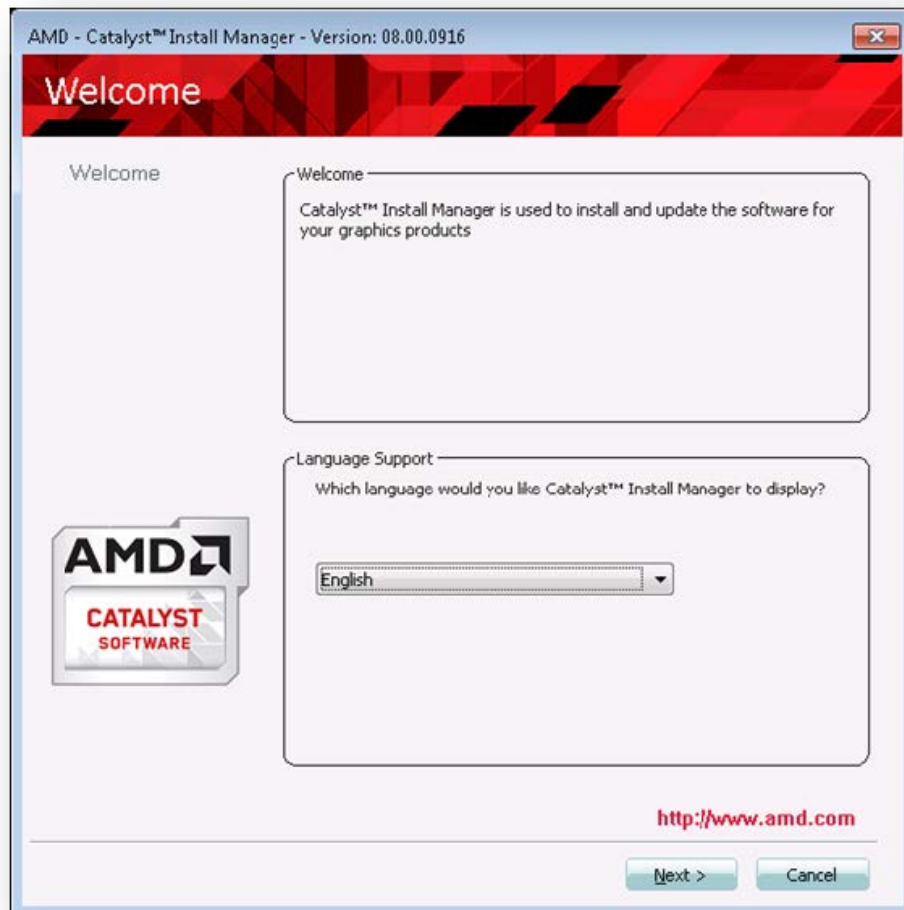
The VM is now ready to run the full range of DirectX and OpenGL graphics applications supported by the GPU.

Install the AMD drivers

To enable GPU operation, install AMD drivers into the VM.

1. Start the VM. In the **Resources** pane, right-click on the VM, and click **Start**.
During this boot process, Citrix Hypervisor dynamically allocates a GPU to the VM.
2. Follow the Windows operating system installation screens.
3. After the operating system installation completes, restart the VM.

4. Copy the 32-bit or 64-bit AMD Windows drivers (AMD Catalyst Install Manager) to the VM.
5. Run the AMD Catalyst Install Manager; select your **Destination Folder**, and then click **Install**.



6. Follow the installer steps to install the driver.
7. To complete the installation, restart your VM.
8. After the VM restarts, check that graphics are working correctly. Open the **Windows Device Manager**, expand **Display adapters**, and ensure that the AMD Graphics Adapter does not have any warning symbols.

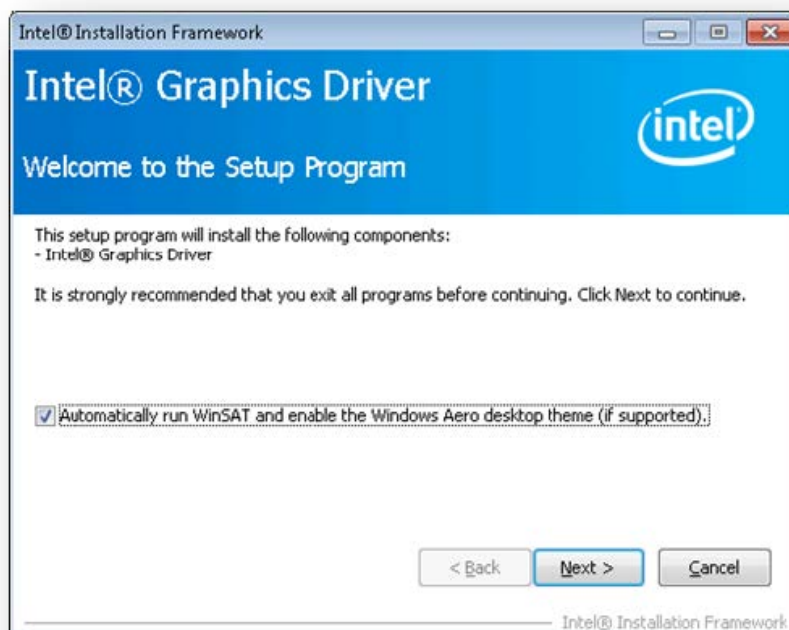
Install the Intel drivers

To enable GPU operation, install Intel drivers into the VM.

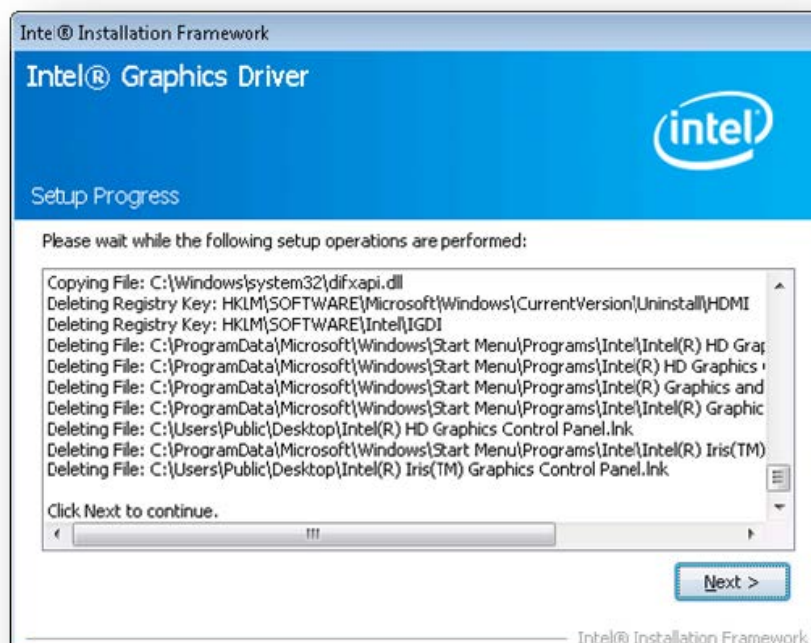
1. Start the **VM**. In the **Resources** pane, right-click on the VM, and click **Start**.

During this boot process, Citrix Hypervisor dynamically allocates a GPU to the VM.

2. Follow the Windows operating system installation screens.
3. After the operating system installation completes, reboot the VM.
4. Copy the 32-bit or 64-bit Intel Windows driver (Intel Graphics Driver) to the VM.
5. Run the **Intel Graphics Driver** setup program
6. Select **Automatically run WinSAT**, and then click **Next**.



7. To accept the License Agreement, click **Yes**, and on the Readme File Information screen, click **Next**.
8. Wait until the setup operations complete. When you are prompted, click **Next**.



9. To complete the installation, you are prompted to restart the VM. Select **Yes**, I want to restart this computer now, and click **Finish**.
10. After the VM restarts, check that graphics are working correctly. Open the Windows Device Manager, expand **Display adapters**, and ensure that the Intel Graphics Adapter does not have any warning symbols.

Note:

You can obtain the latest drivers from the [Intel website](#).

Memory usage

March 18, 2024

Two components contribute to the memory footprint of the Citrix Hypervisor server. 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 Citrix Hypervisor management toolstack (XAPI). Besides providing Citrix Hypervisor 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, Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor server in the following cases:

- You are running many VMs on the server
- 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 server. 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 Citrix Hypervisor 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 command 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 Citrix Hypervisor server, open a local shell and log on as root.
2. Type the following:

```
1 /opt/xensource/libexec/xen-cmdline --set-xen dom0_mem=<nn>M,max:<nn>M
```

Where `<nn>` represents the amount of memory, in MiB, to be allocated to dom0.

3. Restart the Citrix Hypervisor server using XenCenter or the `reboot` command on the Citrix Hypervisor console.

When the host restarts, on the Citrix Hypervisor 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

June 24, 2024

Citrix Hypervisor 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.

Citrix Hypervisor 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 Citrix Hypervisor performance

Customers can monitor the performance of their Citrix Hypervisor servers 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](#).

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.
- Performance metrics are not available for GFS2 SRs and disks on those SRs.

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

Metric Name	Description	Condition	XenCenter Name
<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>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
<code>io_throughput_read_<sr-uuid>-short</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-uuid>-short</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-uuid>-short</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

Metric Name	Description	Condition	XenCenter Name
<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	Free Memory
<code>memory_reclaimed</code>	Host memory reclaimed by squeeze (B).	None	Reclaimed Memory
<code>memory_reclaimed_max</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)

Metric Name	Description	Condition	XenCenter Name
<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)
<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_evicted</code>	Bytes per second evicted from the cache	PVSAccelerator Enabled	PVS-Accelerator eviction rate
<code>pvsaccelerator_read_hits</code>	Reads per second served from the cache	PVSAccelerator Enabled	PVS-Accelerator hit rate
<code>pvsaccelerator_read_misses</code>	Reads per second that cannot be served from the cache	PVSAccelerator Enabled	PVS-Accelerator miss rate
<code>pvsaccelerator_traffic_sent_clients</code>	Bytes per second sent by cached PVS clients	PVSAccelerator Enabled	PVS-Accelerator observed network traffic from clients
<code>pvsaccelerator_traffic_sent_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

Metric Name	Description	Condition	XenCenter Name
pvsaccelerator_space_utilization	Percentage of space used by PVSAccelerator on this host, compared to the total size of the cache storage	PVSAccelerator Enabled	PVS-Accelerator space utilization
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_kilobytes	Memory (KiB) allocation done by the XAPI daemon. Enabled by default.	None	Agent Memory Allocation
xapi_free_memory_kilobytes	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	Citrix Hypervisor High Availability Latency
xapi_live_memory_kilobytes	Live memory (KiB) used by XAPI daemon. Enabled by default.	None	Agent Memory Live
xapi_memory_usage_kilobytes	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>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>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
<code>runstate_partialrun</code>	Fraction of time that some vCPUs are running, and some are blocked	None	vCPUs partial run
<code>runstate_partialcontention</code>	Fraction of time that some vCPUs are runnable and some are blocked	None	vCPUs partial contention

Metric Name	Description	Condition	XenCenter Name
<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
<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

Metric Name	Description	Condition	XenCenter Name
<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`

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 server, 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 may 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

Citrix Hypervisor 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 RRDs use Consolidation Functions instead. The consolidation functions supported by Citrix Hypervisor are:

- AVERAGE
- MIN
- MAX

RRDs exist for individual VMs (including dom0) and the Citrix Hypervisor server. VM RRDs are stored on the host on which they run, or the pool master 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 Citrix Hypervisor RRDs, see the [Citrix Hypervisor Software Development Kit Guide](#).

Analyze RRDs using HTTP

You can download RRDs over HTTP from the Citrix Hypervisor server 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 may not necessarily be unique.

The man page (`rrd2csv --help`) is the definitive help text of the tool.

Alerts

You can configure Citrix Hypervisor to generate alerts based on any of the available Host or VM Metrics. In addition, Citrix Hypervisor provides preconfigured alerts that trigger when hosts undergo certain conditions and states. You can view these alerts using XenCenter or the xe CLI.

View alerts using XenCenter

You can view different types of alerts in XenCenter by clicking **Notifications** and then **Alerts**. The Alerts view displays various types of alerts, including Performance alerts, System alerts, and Software update alerts.

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

1. In the **Resources** pane, select the relevant host, VM, or SR, then click the **General** tab and then **Properties**.
2. Click the **Alerts** tab. You can configure the following alerts:
 - **CPU usage** alerts for a host or VM: Check the **Generate CPU usage alerts** check box, then set the CPU usage and time threshold that trigger the alert
 - **Network usage** alerts for a host or VM: Check the **Generate network usage alerts** check box, then set the network usage and time threshold that trigger the alert.
 - **Memory usage** alerts for a host: Check the **Generate memory usage alerts** check box, and then set the free memory and time threshold that trigger the alert.
 - **Control domain memory usage** alerts for a host: Check the **Generate control domain memory usage alerts** check box, and then set the control domain memory usage and time threshold that trigger the alert.
 - **Disk usage** alerts for a VM: Check the **Generate disk usage alerts** check box, then set the disk usage and time threshold that trigger the alert.
 - **Storage throughput** alerts for an SR: Check the **Generate storage throughput alerts** check box, then set the storage throughput and time threshold that trigger the alert.

Note:

Physical Block Devices (PBD) represent the interface between a specific Citrix Hypervisor server 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 Citrix Hypervisor server alerts, this alert must be configured on the SR.

- 3. 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.
- 4. 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.

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	Citrix Hypervisor 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.

Name	Priority/Severity	Description
ha_pool_overcommitted	3	Potential service loss. High availability is unable to guarantee protection for configured VMs.
ha_poor_drop_in_plan_exists_for3		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.
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 Citrix Hypervisor 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 Master.
pbd_plug_failed_on_server_start	3	The host failed to connect to Storage at boot time.

Name	Priority/Severity	Description
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:** Citrix Hypervisor expects a newer version but can still connect to the current version
- **XenCenter out of date:** XenCenter is too old to connect to Citrix Hypervisor
- **Citrix Hypervisor out of date:** Citrix Hypervisor is an old version that the current XenCenter cannot connect to
- **License expired alert:** Citrix Hypervisor license has expired
- **Missing IQN alert:** Citrix Hypervisor uses iSCSI storage but the host IQN is blank
- **Duplicate IQN alert:** Citrix Hypervisor uses iSCSI storage, and there are duplicate host IQNs

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 Citrix Hypervisor 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`
- `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_xxxxxxxx` (where `xxxxxxxx` 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>
```

Configure email alerts

You can configure Citrix Hypervisor to send email notifications when Citrix Hypervisor servers generate alerts. The mail-alarm utility in Citrix Hypervisor uses sSMTP to send these email notifications. You can enable basic email alerts by using Xencenter or the xe Command Line Interface (CLI). For further configuration of email alerts, you can modify the `mail-alarm.conf` configuration file.

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>
```

Citrix Hypervisor 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.

Further configuration

To further configure the mail-alarm utility in Citrix Hypervisor, create an `/etc/mail-alarm.conf` file containing the following:

```
1 root=postmaster
2 authUser=<username>
3 authPass=<password>
4 mailhub=@MAILHUB@
```

`/etc/mail-alarm.conf` is a user-supplied template for sSMTP's configuration file `ssmtp.conf` and is used for all alerts generated by Citrix Hypervisor servers. It consists of keys where `key=@KEY@` and `@KEY@` is replaced by the corresponding value of `ssmtp-key` in `pool.other_config`. These

values are then passed to `ssmtp`, allowing you to control aspects of the sSMTP configuration using values from `pool.other_config`. Note how `@KEY@` (uppercase) corresponds to `ssmtp-key` (lowercase, prefixed by `ssmtp-`).

For example, if you set the SMTP server:

```
1 xe pool-param-set uuid=pool_uuid other-config:ssmtp-mailhub=smtp.
  example.com
```

and then add the following to your `/etc/mail-alarm.conf` file:

```
1 mailhub=@MAILHUB@
```

`mailhub=@MAILHUB@` becomes `mailhub=smtp.example.com`.

Each SMTP server can differ slightly in its setup and may require extra configuration. To further configure sSMTP, modify its configuration file `ssmtp.conf`. By storing relevant keys in the `mail-alarm.conf` file, you can use the values in `pool.other_config` to configure sSMTP. The following extract from the `ssmtp.conf` man page shows the correct syntax and available options:

```
1 NAME
2     ssmtp.conf - ssmtp configuration file
3
4 DESCRIPTION
5     ssmtp reads configuration data from /etc/ssmtp/ssmtp.conf The file
6     contains keyword-argument pairs, one per line. Lines starting with '#'
7     and empty lines are interpreted as comments.
8
9     The possible keywords and their meanings are as follows (both are case-
10    insensitive):
11
12    Root
13    The user that gets all mail for userids less than 1000. If blank,
14    address rewriting is disabled.
15
16    Mailhub
17    The host to send mail to, in the form host | IP_addr port :
18    <port>. The default port is 25.
19
20    RewriteDomain
21    The domain from which mail seems to come. For user authentication.
22
23    Hostname
24    The full qualified name of the host. If not specified, the host
25    is queried for its hostname.
26
27    FromLineOverride
28    Specifies whether the From header of an email, if any, may over-
29    ride the default domain. The default is "no".
30
```

```
31     UseTLS
32     Specifies whether smtp uses TLS to talk to the SMTP server.
33     The default is "no".
34
35     UseSTARTTLS
36     Specifies whether smtp does a EHLO/STARTTLS before starting
37     TLS
38     negotiation. See RFC 2487.
39
40     TLSCert
41     The file name of an RSA certificate to use for TLS, if required
42     .
43
44     AuthUser
45     The user name to use for SMTP AUTH. The default is blank, in
46     which case SMTP AUTH is not used.
47
48     AuthPass
49     The password to use for SMTP AUTH.
50
51     AuthMethod
52     The authorization method to use. If unset, plain text is used.
53     May also be set to "cram-md5".
```

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

November 14, 2023

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 Citrix Hypervisor 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 Citrix Hypervisor, 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 Citrix Hypervisor 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

Citrix Hypervisor 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%.

Manage virtual machines

February 22, 2024

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](#).

Citrix Hypervisor supports guests with any combination of IPv4 or IPv6 configured addresses.

In Citrix Hypervisor VMs can operate in full virtualized (HVM) 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. Citrix Hypervisor 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. Citrix Hypervisor 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, for example the Demo Linux Virtual Appliance.
- 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 Citrix Hypervisor can be used in newer versions of Citrix Hypervisor. However, templates created in newer versions of Citrix Hypervisor are not compatible with older versions of Citrix Hypervisor. If you created a VM template by using Citrix Hypervisor 8.2, to use it with an earlier version, export the VDIs separately and create the VM again.

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 master 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.

Citrix Hypervisor 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 server 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 server 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 server 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 Citrix Hypervisor server 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 (formerly Citrix VM Tools) 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 Citrix Hypervisor product family software. While installing a Windows operating system, Citrix Hypervisor 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

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

Citrix VM Tools for Linux

Citrix VM Tools for Linux contain a guest agent that provides extra information about the VM to the host.

You must install the Citrix 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, Rocky Linux 8, Rocky Linux 9, or CentOS Stream 9 VMs as these operating systems do not support memory ballooning with the Xen hypervisor.

For more information, see [Install Citrix 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 [Citrix Hypervisor 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.

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 Citrix Hypervisor product family.

Citrix Hypervisor product family virtual device support

The current version of the Citrix Hypervisor 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 Citrix Hypervisor [Hardware Compatibility List](#).

VM block devices Citrix Hypervisor 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.

Windows VMs

April 16, 2024

Installing Windows VMs on the Citrix Hypervisor server requires hardware virtualization support (Intel VT or AMD-V).

Note:

Nested virtualization is not supported for Windows VMs hosted on Citrix Hypervisor.

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 (32-bit)	BIOS	Used to install Windows 10
Windows 10 (64-bit)	BIOS, UEFI, UEFI Secure Boot	Used to install Windows 10 (64-bit)
Windows Server 2016 (64-bit)	BIOS, UEFI, UEFI Secure Boot	Used to install Windows Server 2016 or Windows Server Core 2016 (64-bit)
Windows Server 2019 (64-bit)	BIOS, UEFI, UEFI Secure Boot	Used to install Windows Server 2019 or Windows Server Core 2019 (64-bit)

Template Name	Supported boot modes	Description
Windows Server 2022 (64-bit)	BIOS, UEFI, UEFI Secure Boot	Used to install Windows Server 2022 or Windows Server Core 2022 (64-bit)

Citrix Hypervisor 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 Citrix Hypervisor server, 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.

Guest UEFI boot and Secure Boot

Citrix Hypervisor enables recent versions of Windows guest operating systems to boot in UEFI mode. UEFI boot provides a richer interface for the guest operating systems to interact with the hardware, which can significantly reduce Windows VM boot times.

For these Windows operating systems, Citrix Hypervisor also supports Windows 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.

Note:

Guest UEFI boot was provided as an experimental feature in Citrix Hypervisor 8.0. UEFI-enabled VMs that were created in Citrix Hypervisor 8.0 are not supported in Citrix Hypervisor 8.2. Delete these VMs and create new ones with Citrix Hypervisor 8.2.

Citrix Hypervisor supports UEFI boot and Secure Boot on newly created Windows 10 (64-bit), Windows Server 2016 (64-bit), Windows Server 2019 (64-bit), and Windows Server 2022 (64-bit) VMs. 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 at any time.

Consider the following when enabling UEFI boot on VMs:

- Ensure that the UEFI-enabled VM has at least two vCPUs.
- You can import or export a UEFI-enabled VM created on Citrix Hypervisor 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.
- Use the UEFI settings menu to change the screen resolution of the XenCenter console. For detailed instructions, see [Troubleshooting](#).

Consider the following when enabling UEFI Secure Boot on VMs:

- The Citrix Hypervisor server must be booted in UEFI mode. For more information, see [Network boot installations](#)
- Your resource pool or standalone server must have access to Secure Boot certificates.

Only one Citrix Hypervisor server in the pool requires access to the certificates. When a server joins a pool the certificates on that server are made available to other servers in the pool.

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 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: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 created on a Citrix Hypervisor server that is booted in UEFI

mode and has Secure Boot certificates available, the 'auto' behavior is to enable Secure Boot for the VM. Otherwise, Secure Boot is not enabled.

To create a UEFI-enabled VM from a template supplied with Citrix Hypervisor, 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 Windows installation. If you do not start the Windows 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, Windows debugging cannot be enabled on a VM that is in Secure Boot user mode. To disable Secure Boot, change the VM into Secure Boot setup mode. On your Citrix Hypervisor server, run the following command:

```
1  varstore-sb-state <VM_UUID> setup
```

Keys

UEFI-enabled 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.

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 on Windows VMs](#).

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 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 Citrix Hypervisor 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 HVM VMs from Reseller Option Kit \(BIOS-locked Media\)](#). The option to set user-defined BIOS strings are not available for HVM VMs.

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 Citrix Hypervisor server

Citrix Hypervisor 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.

5. Choose a boot mode for the VM. By default, XenCenter select the most secure boot mode available for the VM operating system version.

Note:

- 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.

6. Select a home server for the VM.

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, Citrix Hypervisor attempts to start the VM on that server. If this action is not possible, an alternate server within the same pool is selected automatically. To choose a home server, click **Place the VM on this server** and select a server 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 server for the VM by analyzing Citrix Hypervisor 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 server 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 server 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 server with the necessary resources.

Click **Next** to continue.

7. Allocate processor and memory resources for the VM. For a Windows 10 VM, the default is 1 virtual CPU and 2,048 MB of RAM. You can also choose to modify the defaults. Click **Next** to continue.
8. 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.
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.

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

13. Follow the OS installation screens and make your selections.

14. After the OS installation completes and the VM reboots, install the 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:

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. Create an ISO Storage Repository:

```
1 xe-mount-iso-sr path_to_iso_sr
```

3. List all of the available ISOs:

```
1 xe cd-list
```

4. 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
```

5. 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.

6. On the XenCenter **Resources** pane, select the VM, and then click the **Console** tab to see the VM console.
7. Follow the OS installation screens and make your selections.
8. 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 (formerly Citrix VM Tools) 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](#).

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.

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 [XenServer VM Tools for Windows](#).

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 [Citrix Hypervisor downloads page](#).
 - a) Log in to your Citrix account to access the downloadable files.
 - b) Expand the product sections on the Citrix Hypervisor downloads page and click into any supported version of Citrix Hypervisor.
 - c) On the product page, go to the **Citrix Hypervisor Main Components** section.

The XenServer VM Tools for Windows are available in a 32-bit and a 64-bit version.
 - d) Download the MSI file and verify your download against the provided SHA256 value.
3. Copy the file to your Windows VM or to a shared drive that the Windows VM can access.
4. Run the `managementagentxXX.msi` file to begin XenServer VM Tools installation.

```
1 Msiexec.exe /package managementagentxXX.msi
```

5. 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.
6. Restart the VM when prompted to complete the installation process.

Update Windows operating systems

This section discusses updating Windows VMs with updated operating systems.

Upgrades to VMs are typically required when moving to a newer version of Citrix Hypervisor. Note the following limitations when upgrading your VMs to a newer version of Citrix Hypervisor:

- Before migrating Windows VMs using live migration, you must upgrade the XenServer VM Tools for Windows on each VM.
- Suspend/Resume operation is not supported on Windows VMs until the XenServer VM Tools for Windows are upgraded.
- The use of certain antivirus and firewall applications can crash Windows VMs, unless the XenServer VM Tools for Windows are upgraded.

We recommend that you do not remove the XenServer VM Tools from your Windows VM before automatically updating the version of Windows on the VM.

Use Windows Update to upgrade the version of the Windows operating system on your Windows VMs.

Note:

Windows installation disks typically provide an upgrade option if you boot them on a server which has an earlier version of Windows already installed. However, if you use Windows Update to update your XenServer VM Tools, do not upgrade the Windows operating system from an installation disk. Instead, use Windows Update.

For information about upgrading the version of the XenServer VM Tools for Windows, see [XenServer VM Tools for 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.

Note:

For older versions of Windows, the [sysprep](#) binaries are on the Windows product CDs in the `\support\tools\deploy.cab` file. These binaries must be copied to your Windows VM before using.

To clone Windows VMs:

1. Create, install, and configure the Windows VM as desired.
2. Apply all relevant Service Packs and updates.
3. Install the latest version of XenServer VM Tools for Windows.
4. Install any applications and perform any other configuration.
5. Run [sysprep](#). This utility shuts down the VM when it completes.
6. Using XenCenter convert the VM into a template.
7. Clone the newly created template into new VMs as required.
8. 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

Note:

Do not restart the original, sys-prepped 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.

For more information about using [sysprep](#), visit the following Microsoft website:

- [The Windows Automated Installation Kit \(AIK\)](#)

Windows VM release notes

There are many versions and variations of Windows with different levels of support for the features provided by Citrix Hypervisor. 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 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

July 15, 2024

XenServer VM Tools for Windows (formerly Citrix VM 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.

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 Citrix Hypervisor. 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 [Updates to XenServer VM Tools for Windows](#) or 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.

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 the Management Agent or update the drivers manually. 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 [Citrix Hypervisor downloads page](#).
 - a) Log in to your Citrix account to access the downloadable files.
 - b) Expand the product sections on the Citrix Hypervisor downloads page and click into any supported version of Citrix Hypervisor.

- c) On the product page, go to the **Citrix Hypervisor Main Components** section.

The XenServer VM Tools for Windows are available in a 32-bit and a 64-bit version.

- d) Download the MSI file and verify your download against the provided SHA256 value.
3. Copy the file to your Windows VM or to a shared drive that the Windows VM can access.
4. Run the `managementagentxXX.msi` file to begin XenServer VM Tools installation.

```
1 Msiexec.exe /package managementagentxXX.msi
```

5. 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. The **Citrix Hypervisor Windows Management Agent Setup** wizard displays the recommended settings. By default, the wizard displays the following settings:
 - Install I/O Drivers Now
 - Allow automatic management agent updates
 - Disallow automatic I/O drivers updates by the management agent
 - Send anonymous usage information to Citrix

If you do not want to allow the automatic updating of the Management Agent, select **Disallow automatic management agent updates** from the list.

If you would like to allow the Management Agent to update the I/O drivers automatically, select **Allow automatic I/O driver updates by the management agent**. However, we recommend that you use Windows Update to update the I/O drivers, not the Management Agent.

Note:

If you have chosen to receive I/O driver updates through the Windows Update mechanism, do not allow the Management Agent to update the I/O drivers automatically.

If you do not want to share anonymous usage information with Citrix, clear the **Send anonymous usage information to Citrix** check box. The information transmitted to Citrix contains the UUID of the VM requesting the update. No other information relating to the VM is collected or transmitted to Citrix.

- Click **Next** and then **Install** to begin the XenServer VM Tools for Windows installation process.
6. Restart the VM when prompted to complete the installation process.

Note:

The XenServer VM Tools for Windows can request to restart with `/quiet /norestart` or `/quiet /forcerestart` specified after the VM has already been restarted once as part of the installation.

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 to install the Management Agent, and to maintain a supported configuration.

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.

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 managementagentx86.msi /quiet /norestart
2 Msiexec.exe /package managementagentx64.msi /quiet /norestart
```

A non-interactive, but non-silent installation can be obtained by running:

```
1 Msiexec.exe managementagentx86.msi /passive
2 Msiexec.exe managementagentx64.msi /passive
```

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
ALLOWDRIVERINSTALL	YES or NO	YES	Install the I/O Drivers now
ALLOWDRIVERUPDATE	YES or NO	NO	Allow the automatic management agent updates to install updated drivers

Parameter	Allowed values	Default	Description
IDENTIFYAUTOUPDATE	YES or NO	YES	Send anonymous usage information to Citrix

For example, to do a silent install of the tools that does not allow future automatic management agent updates and does not send anonymous information to Citrix, run one of the following commands:

```
1 Msiexec.exe /package managementagentx86.msi ALLOWAUTOUPDATE=NO
  IDENTIFYAUTOUPDATE=NO /quiet /norestart
2 Msiexec.exe /package managementagentx64.msi ALLOWAUTOUPDATE=NO
  IDENTIFYAUTOUPDATE=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.

The XenServer VM Tools for Windows are installed by default in the `C:\Program Files\Citrix\XenTools` directory on the VM.

Notes:

- 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 `/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 system reboots after the entire tools installation is complete. 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.

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 particu-

lar adapter might have to be reconfigured following XenServer VM Tools for Windows installation or upgrade.

Update XenServer VM Tools for Windows

Citrix Hypervisor has a simpler mechanism to update I/O drivers (PV drivers) and the Management Agent automatically for Windows VMs. This mechanism enables customers to install updates as they become available.

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 installing or updating the XenServer VM Tools.

Important:

If you are currently using the 8.2.x.x drivers or earlier and want to use the Management Agent MSI file to update to the latest version of the drivers, you must use Device Manager to uninstall the 8.2.x.x drivers from your VM before installing these drivers. If you do not complete this step, the MSI install process fails.

We recommend using the following settings for updating the different components of the XenServer VM Tools for Windows:

1. Set the value of the following registry key to a REG_DWORD value of '3': `HLKM\System\CurrentControlSet\services\xenbus_monitor\Parameters\Autoreboot`
2. Ensure that your VM is configured to receive I/O drivers from Windows Update.
3. Configure the Management Agent to automatically update itself.

The **Virtualization state** section on a VM's **General** tab in XenCenter specifies whether the VM can receive updates from Windows Update. The mechanism to receive I/O driver updates from Windows Update is turned on by default. 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.

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, this might result in unexpected behavior.

The following sections contain information about automatically updating the I/O drivers and the Management Agent.

Update the I/O drivers

You can get I/O driver updates automatically from Microsoft Windows Update, provided:

- You are running Citrix Hypervisor 8.2 Premium Edition, or have access to Citrix Hypervisor through Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.
- You have created a Windows VM using XenCenter issued with Citrix Hypervisor 8.2
- Windows Update is enabled within the VM
- The VM has access to the internet, or it can connect to a WSUS proxy server

Note:

Windows Server Core does not support using Windows Update to install or update the I/O drivers. Instead use the XenServer VM Tools for Windows installer available from the [Citrix Hypervisor downloads page](#). The XenServer VM Tools are available as a downloadable component under all supported versions of Citrix Hypervisor. You must log in to your Citrix account to access these downloads.

Customers can also receive I/O driver updates automatically through the automatic Management Agent update mechanism. You can configure this setting during XenServer VM Tools for Windows installation. For more information, see [Install XenServer VM Tools for Windows](#).

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.

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: `HLKM\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: `HLKM\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.

Update the Management Agent

Citrix Hypervisor enables you to update the Management Agent automatically on both new and existing Windows VMs. By default, Citrix Hypervisor allows the automatic updating of the Management Agent. However, it does not allow the Management Agent to update the I/O drivers automatically. 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:

- You are running Citrix Hypervisor 8.2 Premium Edition, or have access to Citrix Hypervisor through Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement.
- You have installed XenServer VM Tools for Windows issued with Citrix Hypervisor 7.0 or higher
- The Windows VM has access to the Internet

Find the Management Agent version To find out the version of the Management Agent installed on the VM:

1. Navigate to `C:\Program Files\Citrix\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.

Manage Automatic Updates by using the CLI

Citrix Hypervisor 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, automated updates are turned off automatically when the Citrix Virtual Desktops VDA is present and it reports that the machine is non-persistent.

Argument	Values	Description
ALLOWAUTOUPDATE	YES/NO	Allow/disallow auto updating of the Management Agent
ALLOWDRIVERINSTALL	YES/NO	Allow/disallow the XenServer VM Tools for Windows installer to install I/O drivers
ALLOWDRIVERUPDATE	YES/NO	Allow/disallow the Management Agent to update the I/O drivers automatically
IDENTIFYAUTOUPDATE	YES/NO	Allow/disallow the auto update mechanism to send anonymous usage information to Citrix

For example:

```
1 msixexec.exe /i managementagentx64.msi ALLOWAUTOUPDATE=YES  
  ALLOWDRIVERINSTALL=NO \  
2   ALLOWDRIVERUPDATE=NO IDENTIFYAUTOUPDATE=YES
```

Redirect the Management Agent updates

Citrix Hypervisor 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.

To redirect the Management Agent updates:

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 9.2.1.35 and later use <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. Upload the MSI files to an internal web server that your VMs can access.

4. Update the updates file to point to the MSI files on the internal web server.
5. 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\Citrix\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
```

Disable the 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\Citrix\XenTools /t REG_DWORD /v  
   DisableAutoUpdate /d 1
```

To disable 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_enabled=  
   false
```

Modify the automatic I/O driver update settings

During the XenServer VM Tools for Windows installation, you can specify whether you would like to allow the Management Agent to update the I/O drivers automatically. If you prefer to update this setting after completing the XenServer VM Tools for Windows installation process, perform the following steps:

1. On the VM, open a command prompt as an administrator.

2. Run the following command:

```
1 reg.exe ADD HKLM\SOFTWARE\Citrix\XenTools\AutoUpdate /t REG_SZ /v  
2     InstallDrivers /d YES/NO
```

To send anonymous usage information to Citrix:

During the XenServer VM Tools for Windows installation, you can specify whether you would like to send anonymous usage information to Citrix. If you would like to update this setting after completing the XenServer VM Tools for Windows installation process, perform the following steps:

1. On the VM, open a command prompt as an administrator.
2. Run the following command:

```
1 reg.exe ADD HKLM\SOFTWARE\Citrix\XenTools\AutoUpdate REG_SZ /v \  
2     IDENTIFYAUTOUPDATE /d YES/NO
```

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](#).

Uninstall XenServer VM Tools for Windows

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 **Citrix Hypervisor PV 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 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 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`.
- Other:
 - Deletes old entires in **Add or Remove Programs**. This action is the same as that performed 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.

What's new

The version of the XenServer VM Tools for Windows is updated independently of the version of Citrix Hypervisor. 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.

XenServer VM Tools for Windows 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

Improvements in 9.3.1 This release includes the following improvements:

- Improvements to the `uninstall.exe` utility, including the `purge` parameter. For more information, see Uninstall XenServer VM Tools for Windows.
- General improvements to the XenServer VM Tools installer.
- General improvements to string handling of registry keys.

Fixed issues in 9.3.1 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 could 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 shutdown and prevent the upgrade from succeeding.
- On a Windows VM, you might see both the previous and the latest version of the tools or management agent listed in your Installed Programs.
 - (PREVIOUS) Citrix XenServer Windows Management Agent
 - (LATEST) Citrix Hypervisor PV Tools.

Earlier releases

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 sync 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 masters than on other pool members.
- On Windows VM 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 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 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

July 10, 2024

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 Citrix Hypervisor 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 Citrix Hypervisor, 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 Citrix VM Tools for Linux immediately after installing the operating system. For more information, see [Install Citrix 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 Citrix 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.

Citrix Hypervisor 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. Citrix Hypervisor 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. 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 Citrix Hypervisor server:

```
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
```

3. Insert the operating system installation CD into the CD drive on the Citrix Hypervisor server.
4. Open a console to the VM with XenCenter or an SSH terminal and follow the steps to perform the OS installation.
5. Start the VM. It boots straight into the operating system installer:

```
1 xe vm-start uuid=UUID
```

6. Install the guest utilities and configure graphical display. For more information, see [Install the Citrix 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 Citrix Hypervisor 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 HVM VMs from Reseller Option Kit \(BIOS-locked Media\)](#). The option to set user-defined BIOS strings are not available for HVM VMs.

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 Citrix Hypervisor server

Citrix Hypervisor 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. Select a home server for the VM.

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, Citrix Hypervisor attempts to start the VM on that server. If this action is not possible, an alternate server within the same pool is selected automatically. To choose a home server, click **Place the VM on this server** and select a server 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 server for the VM by analyzing Citrix Hypervisor 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 server 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 server 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 server with the necessary resources.

Click **Next** to continue.

6. Allocate processor and memory resources for the VM. Click **Next** to continue.

7. 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.

8. 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**.

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. 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**.

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 Citrix 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 Citrix 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](#)
- RedHat: [Starting a Kickstart installation automatically using PXE](#)
- CentOS: [PXE Setup](#)
- SLES: [Preparing Network Boot Environment](#)

Install Citrix VM Tools for Linux

Although all supported Linux distributions are natively paravirtualized (and don't need special drivers for full performance), Citrix 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, Rocky Linux 8, Rocky Linux 9, or CentOS Stream 9 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 Citrix Hypervisor server. 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 Citrix VM Tools for Linux:

1. Download the Citrix VM Tools for Linux file from the [Citrix Hypervisor 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.

Installation notes for Linux distributions

This following 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](#).

Red Hat Enterprise Linux* 7 (32-/64-bit)

The new template for these guests specifies 2 GB RAM. This amount of RAM is a requirement for a successful install of v7.4 and later. For v7.0 - v7.3, the template specifies 2 GB RAM, but as with previous versions of Citrix Hypervisor, 1 GB RAM is sufficient.

Note:

This information applies to both Red Hat and Red Hat derivatives.

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, Citrix Hypervisor 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 Citrix VM Tools for Linux (see [Install the Citrix 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.

Upgrade from PV to HVM guests

PV mode guests are not supported in Citrix Hypervisor 8.2. Before upgrading your server to Citrix Hypervisor 8.2, follow these steps to upgrade your existing PV Linux guests to supported versions.

1. Upgrade the guest operating system to a version that is supported by Citrix Hypervisor. Perform the upgrade by using the mechanism provided by that Linux distribution.

After this step, the upgraded guest is still PV mode, which is not supported and has known issues.

2. Use the `pv2hvm` script to convert the newly upgraded guest to the supported HVM mode.

On the Citrix Hypervisor server, open a local shell, log on as root, and enter the following command:

```
1 /opt/xensource/bin/pv2hvm vm_name
```

Or

```
1 /opt/xensource/bin/pv2hvm vm_uuid
```

3. Restart the VM to complete the process.

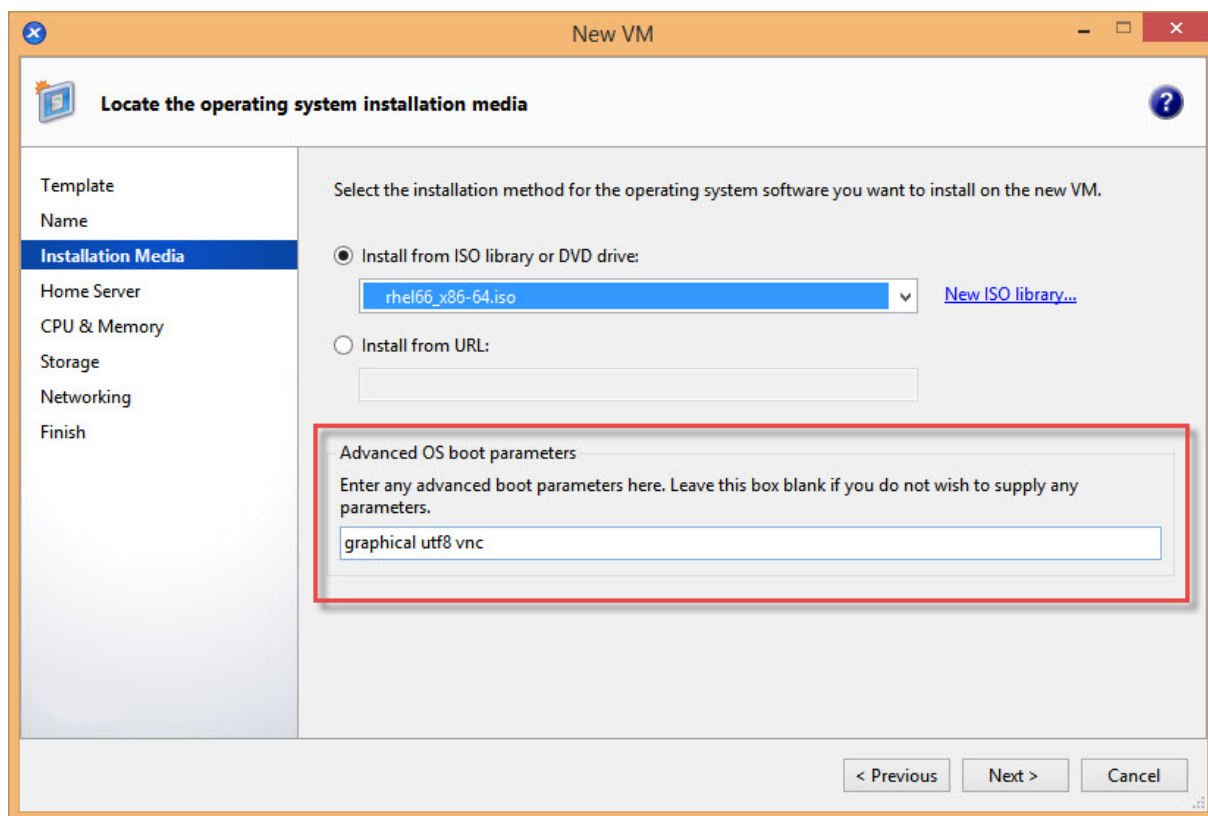
Linux VM release notes

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 7

After migrating or suspending the VM, RHEL 7 guests might freeze during resume. For more information, see Red Hat issue [1141249](#).

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, Rocky Linux 8, Rocky Linux 9, or CentOS Stream 9 VMs as these operating systems do not support memory ballooning with the Xen hypervisor.

Red Hat Enterprise Linux 9

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

CentOS 7

For the list of CentOS 7 release notes, see [Red Hat Enterprise Linux 7](#).

CentOS Stream 9

If you attempt to shut down a CentOS Stream 9 VM by using XenCenter or the xe CLI, the shutdown process pauses and times out after 1200s. This behaviour is caused by a kernel issue in kernel-5.14.0-362.el9.

The following workarounds are available:

- To work around a single instance of the issue, you can shut down the VM from inside the guest operating system.
- To prevent the issue from occurring for your VM, downgrade your VM to use kernel-5.14.0-354.el9 by running the following commands in the VM:

```
1 # yum install kernel-5.14.0-354.el9
2 # grubby --set-default /boot/vmlinuz-5.14.0-354.el9.x86_64
```

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

Rocky 9 and CentOS Stream 9 VMs require that the host has a CPU compatible with x86-64-v2 instruction set or higher. For more information, see <https://access.redhat.com/solutions/6833751>.

Oracle Linux 7

For the list of Oracle Linux 7 release notes, see [Red Hat Enterprise Linux 7](#).

Scientific Linux 7

For the list of Scientific Linux 7 release notes, see [Red Hat Enterprise Linux 7](#).

Rocky Linux 9

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

Rocky 9 and CentOS Stream 9 VMs require that the host has a CPU compatible with x86-64-v2 instruction set or higher. For more information, see <https://access.redhat.com/solutions/6833751>.

Debian 10

If you install Debian 10 (Buster) by using PXE network boot, do not add `console=tty0` to the boot parameters. This parameter can cause issues with the installation process. Use only `console=hvc0` in the boot parameters.

For more information, see Debian issues [944106](#) and [944125](#).

Debian 11

When installing Debian 11 32-bit on a VM using a QEMU emulated network device, the installation might fail. This issue is caused by the Xen PV drivers being missing from the installer kernel.. For more information, see <https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=818481>.

This issue is not present in the 64-bit Debian 11.

Debian 12

Debian 12 (Bookworm) does not support GPU passthrough or vGPUs.

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 a 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

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.

We recommend to work around this issue by disabling Wayland or upgrading your Mesa drivers:

- 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.

- Update your Mesa drivers

Run the following commands:

```
1 sudo add-apt-repository ppa:kisak/kisak-mesa
2 sudo apt update
3 sudo apt upgrade
```

Restart the VM to complete the process.

Migrate VMs

July 15, 2024

You can migrate running VMs by using *live migration* and *storage live migration* and move a VMs 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 Citrix Hypervisor. This feature enables you to move a running VM from one host to another host, when the VMs disks 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.

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.

Intel GVT-g is not compatible with live migration, storage live migration, or VM Suspend. For information, see [Graphics](#).

Storage live migration

Notes:

- Do not use storage live migration in Citrix Virtual Desktops deployments.
- Storage live migration cannot be used on VMs that have changed block tracking enabled. Disable changed block tracking before attempting storage live migration.
- Storage live migration cannot be used on VMs whose VDIs are on a GFS2 SR.

Storage live migration also allows VMs to be moved from one host to another, where the VMs 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 Citrix Hypervisor pools (for example from a development environment to a production environment).
- Upgrade and update standalone Citrix Hypervisor servers without any VM downtime.
- Upgrade Citrix Hypervisor 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.
- You cannot migrate a VM from a source pool that does *not* have hotfix [XS82ECU1033](#) installed, to a destination pool that *does* and has port 80 closed. To do so, install hotfix [XS82ECU1033](#) on the source pool or temporarily reopen port 80 of the destination pool. For more information, see [Restrict use of port 80](#).

Compatibility requirements

When migrating a VM with live migration or storage live migration, VM and the target host must meet the following compatibility requirements for the migration to proceed:

- The target host must have the same or a more recent version of Citrix Hypervisor installed as the source host.
- XenServer VM Tools for Windows must be installed on each Windows VM that you want to migrate.
- Storage live migration only: If the CPUs on the source and target host are different, the target CPU must provide at least the entire feature set as the source CPU. So, it is unlikely to be possible to move a VM between, for example, AMD and Intel processors.
- VMs with checkpoint cannot be migrated.
- Storage live migration only: VMs with more than six attached VDIs cannot be migrated.
- 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.
- Storage live 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:

- Storage live migration cannot be used with VMs created by Machine Creation Services.
- VMs using SR-IOV cannot be migrated. For more information, see [Use SR-IOV enabled NICs](#)
- VM performance is reduced during migration.
- If using the high availability feature, ensure the VM being migrated is not marked as protected.
- 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.
- Intel GVT-g is not compatible with live migration and storage live migration. For more information, see [Graphics overview](#)
- VMs that have the `on-boot` option set to `reset` cannot be migrated. For more information, see [Intellicache](#).

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. From the **Destination** list, select a standalone server or a pool.
3. From the **Home Server** list, select a server to assign as the home server for the VM and click **Next**.
4. In the **Storage** tab, 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.
5. From the **Storage network** list, select a network on the destination pool that is 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. Review the configuration settings and click **Finish** to start migrating the VM.

If you are migrating from an older version of XenServer or Citrix Hypervisor, 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

- Do not use storage live migration in Citrix Virtual Desktops deployments.
- 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

July 15, 2024

Citrix Hypervisor allows you to import VMs from and export them to various different formats. Using the XenCenter Import wizard, you can import VMs from disk images (VHD and VMDK), Open Virtualization Format (OVF and OVA) and Citrix Hypervisor 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 Citrix Hypervisor. 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 Citrix Hypervisor XVA format.

You can also use the xe CLI to import VMs from and export them to Citrix Hypervisor 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	Description
Open Virtualization Format (OVF and OVA)	OVF is an open standard for packaging and distributing a virtual appliance consisting of one or more VMs.
Disk image formats (VHD and VMDK)	Virtual Hard Disk (VHD) and Virtual Machine Disk (VMDK) format disk image files can be imported using the Import wizard. Importing a disk image may be appropriate when there is a virtual disk image available, with no OVF metadata associated.
Citrix Hypervisor XVA format	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 Citrix Hypervisor vApps and VMs with other virtualization platforms that support OVF
- Save more than one VM
- Secure a vApp or VM from corruption and tampering
- Include a license agreement
- Simplify vApp distribution by storing an OVF package in an OVA file

Consider using XVA format to:

- Import and export VMs from a script with a 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. Citrix Hypervisor 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 Citrix Hypervisor server 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 and VMDK)

Using XenCenter, you can import disk images in the Virtual Hard Disk (VHD) and Virtual Machine Disk (VMDK) formats. Exporting standalone disk images is not supported.

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 the Import wizard requires to recreate a VM from its disk image. This information includes the number of disk images associated with the VM, the processor, storage, network, memory requirements and so on. Without this information, it can be much more complex and error-prone to recreate the VM.

XVA format

XVA is a virtual appliance format specific to Citrix Hypervisor, 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 Citrix Hypervisor server and then imported into another Citrix Hypervisor server 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 Citrix Hypervisor, you might have to configure the VM before it boots properly on the Citrix Hypervisor server.

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 Citrix Hypervisor. 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 Citrix Hypervisor 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 Advanced Options page of the XenCenter Import wizard. Specify a location where the Fixup ISO is copied so that Citrix Hypervisor can use it.

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 a Windows system is detected, a generic critical boot device driver is extracted from the driver database of the installed OS and registered with the OS. This process is especially important for older Windows operating systems when the boot device is changed between a SCSI and IDE interface.

If certain virtualization tool sets are discovered in the VM, they are disabled to prevent performance problems and unnecessary event messages.

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.

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 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. 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, and then click **Next** to continue.
3. Review and accept EULAs, if applicable.

If the package you are importing includes any EULAs, accept them and click **Next** to continue. When no EULAs are included in the package, the wizard skips this step and advances straight to the next page.

4. Specify the pool or host to which you want to import the VMs. To select a host or pool, choose from the **Import VM(s) to** list.

Click **Next** to continue.

5. Configure storage for the imported VMs: Choose one or more storage repositories on which to place the imported virtual disks, and then click **Next** to continue.

To place all the imported virtual disks on the same SR, select **Place all imported VMs on this target SR**. Select an SR from the list.

To place the virtual disks of incoming VMs onto different SRs, select **Place imported VMs on the specified target SRs**. For each VM, select the target SR from the list in the SR column.

6. Configure networking for the imported VMs: map the virtual network interfaces in the VMs you are importing to target networks 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.
7. Specify security settings: If the selected OVF/OVA package is configured with security features, such as certificates or a manifest, specify the information necessary, and then click **Next** to continue.

Different options appear on the Security page depending on which security features have been configured on the OVF appliance:

- If the appliance is signed, a **Verify digital signature** check box appears, automatically selected. 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. Clear the **Verify digital signature** check box if you do not want to verify the signature.
- If the appliance includes a manifest, a **Verify manifest content** check box appears. Select this check box to have the wizard verify the list of files in the package.

When packages are digitally signed, the associated manifest is verified automatically, so the **Verify manifest content** check box does not appear on the Security page.

Note:

VMware Workstation 7.1.x OVF files fail to import when you choose to verify the manifest. This failure occurs because VMware Workstation 7.1.x produces an OVF file with a manifest

that has invalid SHA-1 hashes. If you do not choose to verify the manifest, the import is successful.

8. Enable Operating System Fixup: If the VMs in the package you are importing were built on a virtualization platform other than Citrix Hypervisor, select the **Use Operating System Fixup** check box. Select an ISO SR where the Fixup ISO can be copied so that Citrix Hypervisor can access it. For more information about this feature, see [Operating system fixup](#).

Click **Next** to continue.

9. 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 **Logs** tab. When the newly imported VM is available, it appears in the **Resources** pane, and the new vApp appears in the **Manage vApps** dialog box.

Note:

After using XenCenter to import an OVF package that contains Windows operating systems, you must set the `platform` parameter.

1. Set the `platform` parameter to `device_id=0002`. For example:

```
1 xe vm-param-set uuid=VM uuid platform:device_id=0002
```

2. Set the `platform` parameter to `viridian=true`. For example:

```
1 xe vm-param-set uuid=VM uuid platform:viridian=true
```

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 Citrix Hypervisor 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. 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, and then click **Next** to continue.
3. Specify the VM name and allocate CPU and memory resources.

On the **VM Definition** page, enter the name of the new VM to be created from the imported disk image, and allocate CPU and initial memory resources. All of these values can be adjusted later, after the new VM has been created. For more information, see [VM CPU and Memory Allocation](#). 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.

4. Specify the pool or host to which you want to import the VMs. To select a host or pool, choose from the **Import VM(s) to** list.

Click **Next** to continue.

5. Configure storage for the imported VMs: Select one or more storage repositories on which to place the imported virtual disks, and then click **Next** to continue.

To place all the imported virtual disks on the same SR, select **Place all imported VMs on this target SR**. Select an SR from the list.

To place the virtual disks of incoming VMs onto different SRs, select **Place imported VMs on the specified target SRs**. For each VM, select the target SR from the list in the SR column.

6. Configure networking for the imported VMs: map the virtual network interfaces in the VM you are importing to target networks 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.

7. **Specify the boot option.**

On the **Boot options** page, specify the boot mode for the new VM. Select **BIOS Boot** to boot the VM in legacy BIOS mode.

Citrix Hypervisor supports UEFI boot on certain operating systems. For more information about the guest operating systems that are supported for UEFI boot, see [Guest UEFI boot and Secure Boot](#).

Select **UEFI Boot** to boot the VM in UEFI mode.

8. Enable Operating System Fixup: If the disk images you are importing were built on a virtualization platform other than Citrix Hypervisor, select the **Use Operating System Fixup** check box. Select an ISO SR where the Fixup ISO can be copied so that Citrix Hypervisor can access it. For more information about this feature, see [Operating system fixup](#).

Click **Next** to continue.

9. 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 **Logs** tab. When the newly imported VM is available, it appears in the **Resources** pane.

Note:

After using XenCenter to import a disk image that contains Windows operating systems, you must set the `platform` parameter. The value of this parameter varies according to the version of Windows contained in the disk image:

- For Windows Server 2016 and later, set the `platform` parameter to `device_id=0002`. For example:

```
1 xe vm-param-set uuid=VM uuid platform:device_id=0002
```

- For all other versions of Windows, set the `platform` parameter to `viridian=true`. For example:

```
1 xe vm-param-set uuid=VM uuid platform:viridian=true
```

Import VMs from XVA

Note:

The VM Admin role can only import XVA files into a pool with a shared SR. The VM Admin role does not have permission to import an XVA file into a server or a pool without shared storage.

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 server with a different CPU type. For example, a Windows VM exported from a server with an Intel VT Enabled CPU might not run when imported to a server 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 **Logs** tab. 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 Citrix Hypervisor server, 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 Citrix Hypervisor server, 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 via the xe CLI.

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 **Logs** tab. To cancel an export in progress, click the **Logs** tab, 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 Citrix Hypervisor server, 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 server with a different CPU type. For example, a Windows VM exported from a server with an Intel VT Enabled CPU might not run when imported to a server 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 **Logs** tab. To cancel an export in progress, click the **Logs** tab, 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.

Delete VMs

January 10, 2023

You can delete VMs by using the xe CLI or XenCenter.

Deleting a virtual machine removes its configuration and its filesystem from the server. 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

January 9, 2023

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 servers 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**.

VM memory

July 23, 2024

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 Citrix Hypervisor 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 servers are full, DMC allows you to start more VMs on these servers, reducing the amount of memory allocated to the running VMs proportionally.

What is Dynamic Memory Control (DMC)?

Citrix Hypervisor 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 server.

Without DMC, when a server 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, Citrix Hypervisor attempts to reclaim memory by automatically reducing the current memory allocation of running VMs within their defined memory ranges. Citrix Hypervisor attempts to reclaim memory even when the server 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. Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor to ensure that this amount of memory is allocated to the VM. If new VMs are required to start on “full” servers, 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. Citrix Hypervisor *guarantees* always to assign each VM memory within its specified DMR when using DMC.

The concept of static range

Many operating systems that Citrix Hypervisor supports do not fully ‘understand’ the notion of dynamically adding or deleting memory. As a result, Citrix Hypervisor 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 Citrix Hypervisor. 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 Citrix Hypervisor.

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, Citrix Hypervisor 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. Citrix Hypervisor 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. Citrix Hypervisor 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. Citrix Hypervisor 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 Citrix Hypervisor to repartition host memory dynamically in response to changing numbers of guests, or changing host memory pressure. Citrix Hypervisor 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 Citrix Hypervisor takes care of the rest.

Memory constraints

Citrix Hypervisor allows administrators to use all memory control operations with any guest operating system. However, Citrix Hypervisor enforces the following memory property ordering constraint for all guests:

```
0 < memory-static-min <= memory-dynamic-min <= memory-dynamic-max <= memory-static-max
```

Citrix Hypervisor 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. Citrix Hypervisor 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:

$$0 < \text{memory-}\textbf{static}\text{-min} \leq \text{memory-dynamic-min} \leq \text{memory-dynamic-max} \leq \text{memory-}\textbf{static}\text{-max}$$

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. Citrix Hypervisor 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
```

Demo Linux virtual appliance

January 9, 2023

We provide a fully functional installation of a Demo Linux Virtual Appliance, based on a CentOS 7.5 distribution.

The appliance is available for download, in a single `xva` file from the [Citrix Hypervisor Download](#) page.

The `xva` file can be quickly imported into XenCenter to create a fully working Linux Virtual Machine. No additional configuration steps are required.

The Demo Linux Virtual Appliance enables you to deploy a VM quickly and simply. Use this appliance to test Citrix Hypervisor product features such as live migration and high availability.

The Demo Linux Virtual Appliance comes with the following items already set up:

- Citrix VM Tools for Linux
- Pre-configured networking connectivity
- A web server for test purposes

Warning:

Do not use the Demo Linux Virtual Appliance for running production workloads.

Import the Demo Linux virtual appliance

1. Download the Demo Linux Virtual Appliance from the [Citrix Hypervisor Download](#) page.
Customers require access to **My Account** to access this page. If you do not have an account, you can register on the Citrix home page.
2. In the **Resources** pane, select a host or a Pool, then right-click and select **Import**. The Import Wizard is displayed.
3. Click **Browse** and navigate to the location of the downloaded Demo Linux Virtual Appliance `xva` file on your computer.
4. Click **Next**.
5. Select the target Citrix Hypervisor server or pool, then click **Next**.
6. Select a storage repository on which to create the virtual appliance's disk, then click **Next**.
7. Click **Finish** to import the virtual appliance.

Note:

When you first start the VM, you are prompted to enter a root password. The IP address of the VM is then displayed. Ensure that you record the IP address, as it is useful for test purposes.

Useful tests

This section lists some useful tests to carry out to ensure that your Demo Linux Virtual Appliance is correctly configured.

1. Test that you have external networking connectivity.

Log in to the VM from the XenCenter console. Run this command to send ping packets to Google and back:

```
1 ping -c 10 google.com
```

Other installed networking tools include `ifconfig`, `netstat`, and `tracert`.

2. Using the IP address displayed on VM boot, test that you can ping the VM from an external computer.

3. Test that the web server is configured.

In a web browser, enter the VM IP address. The “Demonstration Linux Virtual Machine” page opens. This page displays simple information about the VM mounted disks, their size, location, and usage.

You can also use the webpage to mount a disk.

Mount a disk using the Demonstration Linux Virtual Machine webpage

1. In XenCenter, add a virtual disk to your VM. Select the VM in the **Resources** pane, open the **Storage** tab, and then click **Add**.
2. Enter the name of the new virtual disk and, optionally, a description.
3. Enter the size of the new virtual disk.

Ensure that the storage repository where the virtual disk is stored has sufficient space for the new virtual disk.

4. Select the SR where the new virtual disk is stored.
5. Click **Create** to add the new virtual disk and close the dialog box.
6. Click the **Console** tab, and use your normal tools to partition and format the disk as required.
7. Refresh the Demonstration Linux Virtual Machine webpage, the new disk is displayed.
8. Click **Mount**. This action mounts the disk, and filesystem information is displayed.

For more information on adding virtual disks, see the [XenCenter documentation](#).

Advanced notes for virtual machines

January 16, 2024

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:

```
1 xe vdi-param-set uuid=vdi_uuid on-boot=reset
```

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 Citrix Hypervisor servers

To make an ISO library available to Citrix Hypervisor servers, 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 Citrix Hypervisor 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 (formerly Citrix VM Tools). 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 running a reliable NTP service in the control domain and all Windows VMs.

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. Citrix Hypervisor uses the hardware clock setting of the VM to synchronize the VM. Citrix Hypervisor 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 Citrix Hypervisor, 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 Citrix Hypervisor 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 Citrix Hypervisor BIOS strings.

Note:

If a VM doesn't have BIOS strings set when it starts, the standard Citrix Hypervisor BIOS strings are inserted into it and the VM becomes BIOS-generic.

BIOS-customized

For HVM VMs 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 server 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 HVM 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)

Citrix Hypervisor enables you to assign a physical GPU in the Citrix Hypervisor server 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 Citrix Hypervisor 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 Citrix Hypervisor server. 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 Citrix Hypervisor server and then restart the machine. Upon restart, Citrix Hypervisor 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 Citrix Hypervisor server.

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, Citrix Hypervisor 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

Citrix Hypervisor 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/cding_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.

Enable VNC for Linux VMs

January 9, 2023

VMs might not be set up to support Virtual Network Computing (VNC), which Citrix Hypervisor uses to control VMs remotely, by default. Before you can connect with XenCenter, ensure that the VNC server and an X display manager are installed on the VM and properly configured. This section describes

how to configure VNC on each of the supported Linux operating system distributions to allow proper interactions with XenCenter.

For CentOS-based VMs, use the instructions for the Red Hat-based VMs below, as they use the same base code to provide graphical VNC access. CentOS X is based on Red Hat Enterprise Linux X.

Enable a graphical console on Debian VMs

Note:

Before enabling a graphical console on your Debian VM, ensure that you have installed the Citrix VM Tools for Linux. For more information, see [Install the Citrix VM Tools for Linux](#).

The graphical console for Debian virtual machines is provided by a VNC server running inside the VM. In the recommended configuration, a standard display manager controls the console so that a login dialog box is provided.

1. Install your Debian guest with the desktop system packages, or install GDM (the display manager) using apt (following standard procedures).
2. Install the Xvnc server using `apt-get` (or similar):

```
1 apt-get install vnc4server
```

Note:

The Debian Graphical Desktop Environment, which uses the Gnome Display Manager version 3 daemon, can take significant CPU time. Uninstall the Gnome Display Manager `gdm3` package and install the `gdm` package as follows:

```
1 apt-get install gdm
2 apt-get purge gdm3
```

3. Set up a VNC password (not having one is a serious security risk) by using the `vncpasswd` command. Pass in a file name to write the password information to. For example:

```
1 vncpasswd /etc/vncpass
```

4. Modify your `gdm.conf` file (`/etc/gdm/gdm.conf`) to configure a VNC server to manage display 0 by extending the `[servers]` and `[daemon]` sections as follows:

```
1 [servers]
2 0=VNC
3 [daemon]
4 VTAllocation=false
5 [server-VNC]
6 name=VNC
```

```
7    command=/usr/bin/Xvnc -geometry 800x600 -PasswordFile /etc/  
    vncpass BlacklistTimeout=0  
8    flexible=true
```

5. Restart GDM, and then wait for XenCenter to detect the graphical console:

```
1 /etc/init.d/gdm restart
```

Note:

You can check that the VNC server is running using a command like `ps ax | grep vnc`.

Enable a graphical console on Red Hat, CentOS, or Oracle Linux VMs

Note:

Before setting up your Red Hat VMs for VNC, be sure that you have installed the Citrix VM Tools for Linux. For more information, see [Install the Citrix VM Tools for Linux](#).

To configure VNC on Red Hat VMs, modify the GDM configuration. The GDM configuration is held in a file whose location varies depending on the version of Red Hat Linux you are using. Before modifying it, first determine the location of this configuration file. This file is modified in several subsequent procedures in this section.

Determine the location of your VNC configuration file

If you are using Red Hat Linux, the GDM configuration file is `/etc/gdm/custom.conf`. This file is a split configuration file that contains only user-specified values that override the default configuration. This type of file is used by default in newer versions of GDM. It is included in these versions of Red Hat Linux.

Configure GDM to use VNC

1. As root on the text CLI in the VM, run the command `rpm -q vnc-server gdm`. The package names `vnc-server` and `gdm` appear, with their version numbers specified.

The package names that are displayed show the packages that are already installed. If you see a message that says that a package is not installed, you might have not selected the graphical desktop options during installation. Install these packages before you can continue. For details regarding installing more software on your VM, see the appropriate Red Hat Linux x86 Installation Guide.

2. Open the GDM configuration file with your preferred text editor and add the following lines to the file:

```
1 [server-VNC]
2 name=VNC Server
3 command=/usr/bin/Xvnc -SecurityTypes None -geometry 1024x768 -
   depth 16 \
4 -BlacklistTimeout 0
5 flexible=true
```

With configuration files on Red Hat Linux, add these lines into the empty `[servers]` section.

3. Modify the configuration so that the `Xvnc` server is used instead of the standard X server:

- `0=Standard`

Modify it to read:

`0=VNC`

- If you are using Red Hat Linux, add the above line just below the `[servers]` section and before the `[server-VNC]` section.

4. Save and close the file.

Restart GDM for your change in configuration to take effect, by running the command `/usr/sbin/gdm-restart`.

Note:

Red Hat Linux uses runlevel 5 for graphical startup. If your installation starts up in runlevel 3, change this configuration for the display manager to be started and get access to a graphical console. For more information, see [Check Run levels](#).

Firewall settings

The firewall configuration by default does not allow VNC traffic to go through. If you have a firewall between the VM and XenCenter, allow traffic over the port that the VNC connection uses. By default, a VNC server listens for connections from a VNC viewer on TCP port `5900 + n`, where `n` is the display number (usually zero). So a VNC server setup for Display-0 listens on TCP port 5900, Display-1 is TCP-5901, and so on. Consult your firewall documentation to ensure that these ports are open.

If you want to use IP connection tracking or limit the initiation of connections to be from one side only, further configure your firewall.

To configure Red Hat-base VMS firewall to open the VNC port:

1. For Red Hat Linux, use `system-config-securitylevel-tui`.
2. Select **Customize** and add 5900 to the other ports list.

Alternatively, you can disable the firewall until the next reboot by running the command `service iptables stop`, or permanently by running `chkconfig iptables off`. This configuration can expose extra services to the outside world and reduce the overall security of your VM.

VNC screen resolution

After connecting to a VM with the graphical console, the screen resolution sometimes doesn't match. For example, the VM display is too large to fit comfortably in the Graphical Console pane. Control this behavior by setting the VNC server `geometry` parameter as follows:

1. Open the GDM configuration file with your preferred text editor. For more information, see [Determine the Location of your VNC Configuration File](#).
2. Find the `[server-VNC]` section you added above.
3. Edit the command line to read, for example:

```
1 command=/usr/bin/Xvnc -SecurityTypes None -geometry 800x600
```

The value of the `geometry` parameter can be any valid screen width and height.

4. Save and close the file.

Enable VNC for RHEL, CentOS, or OEL VMs

If you are using Red Hat Linux, the GDM configuration file is `/etc/gdm/custom.conf`. This file is a split configuration file that contains only user-specified values that override the default configuration. By default, this type of file is used in newer versions of GDM and is included in these versions of Red Hat Linux.

During the operating system installation, select **Desktop** mode. On the RHEL installation screen, select **Desktop > Customize now** and then click **Next**:

The default installation of Red Hat Enterprise Linux is a basic server install. You can optionally select a different set of software now.

- ☐ Basic Server
- ☐ Database Server
- ☐ Web Server
- ☐ Identity Management Server
- ☐ Virtualization Host
- ☒ Desktop
- ☐ Software Development Workstation
- ☐ Minimal

Please select any additional repositories that you want to use for software installation.

- ☐ High Availability
- ☐ Load Balancer
- ☒ Red Hat Enterprise Linux
- ☐ Resilient Storage

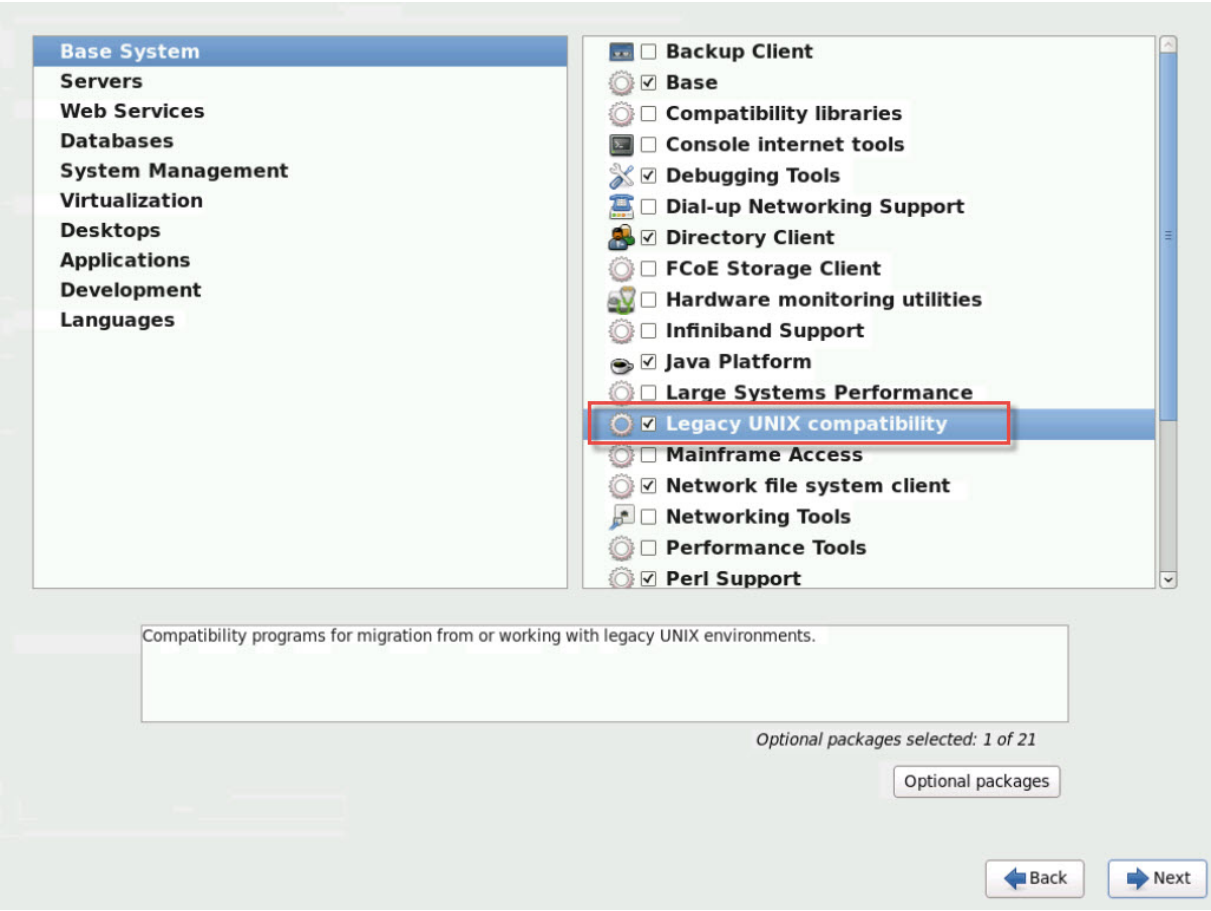
[+ Add additional software repositories](#) [Modify repository](#)

You can further customize the software selection now, or after install via the software management application.

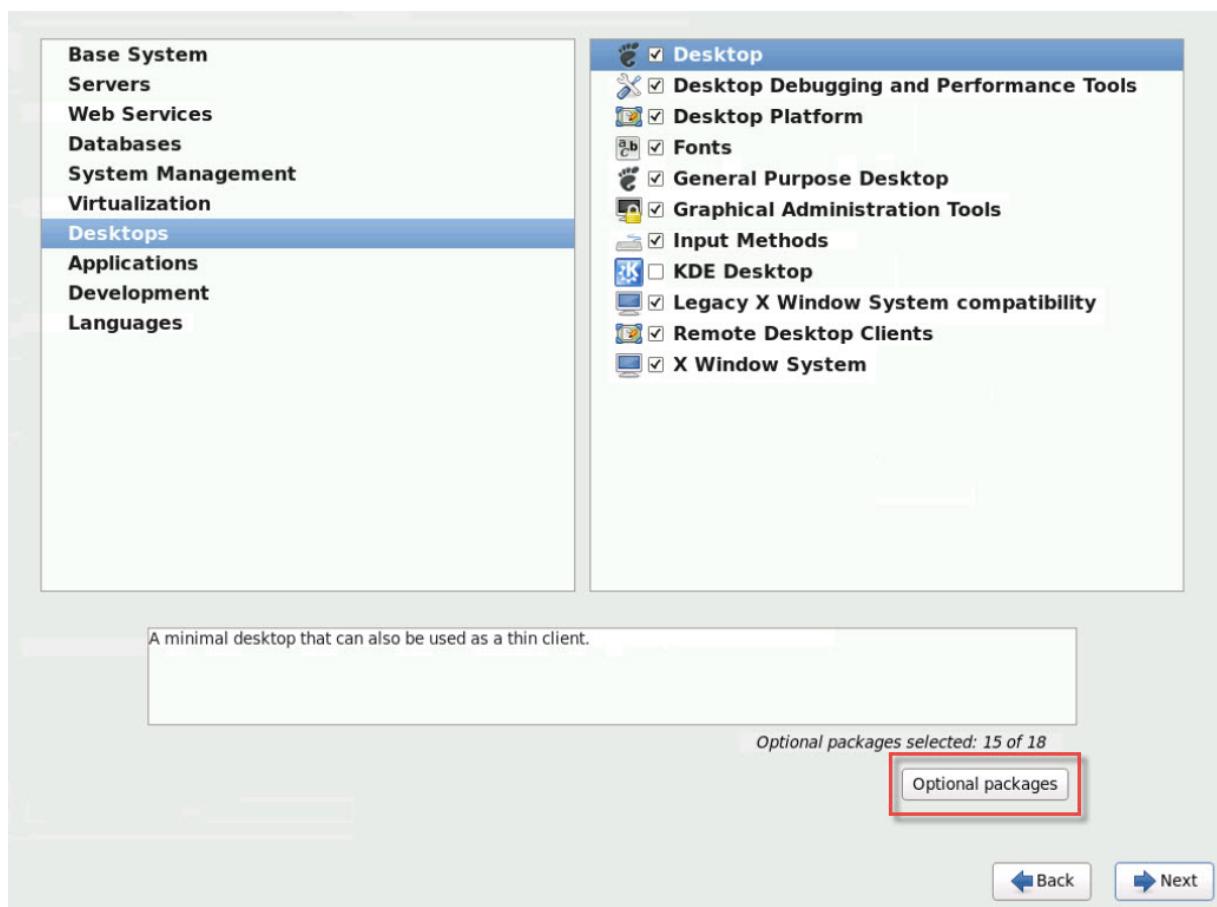
- ☐ Customize later
- ☒ Customize now

[Back](#) [Next](#)

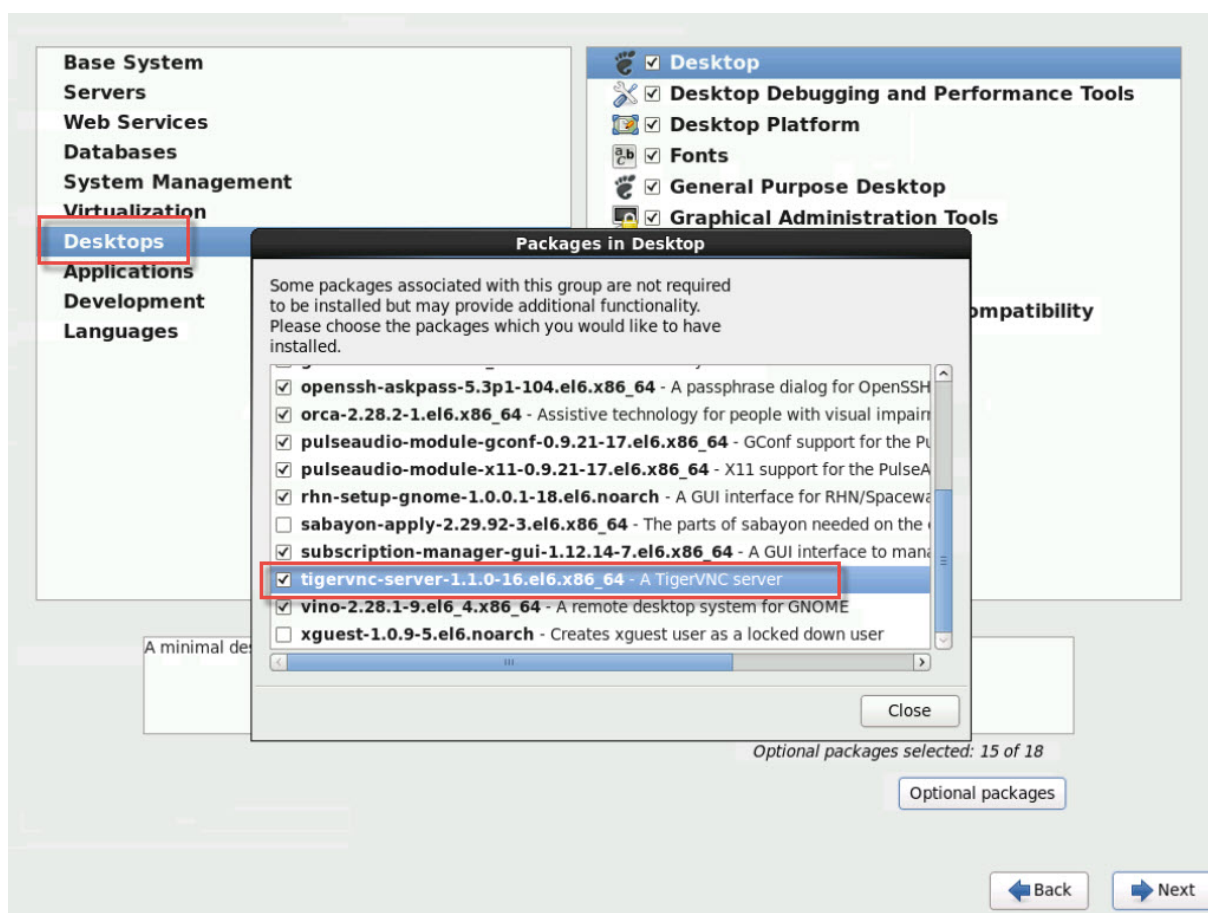
This action displays the Base System screen, ensure that **Legacy UNIX compatibility** is selected:



Select **Desktops > Optional packages**, then click **Next**:



This action displays the **Packages in Desktop** window, select **tigervnc-server-<version_number>** and then click **Next**:



Work through the following steps to continue the setup of your RHEL VMs:

1. Open the GDM configuration file with your preferred text editor and add the following lines to the appropriate sections:

```

1  [security]
2  DisallowTCP=false
3
4  [xdmcp]
5  Enable=true

```

2. Create the file, `/etc/xinetd.d/vnc-server-stream`:

```

1  service vnc-server
2  {
3
4      id = vnc-server
5      disable = no
6      type = UNLISTED
7      port = 5900
8      socket_type = stream
9      wait = no
10     user = nobody
11     group = tty

```

```
12         server = /usr/bin/Xvnc
13         server_args = -inetd -once -query localhost -
14             SecurityTypes None \
15             -geometry 800x600 -depth 16
16     }
```

3. Enter the following command to start the `xinetd` service:

```
1 # service xinetd start
```

4. Open the file `/etc/sysconfig/iptables`. Add the following line above the line reading, `-A INPUT -j REJECT --reject-with icmp-host-prohibited`:

```
1 -A INPUT -m state --state NEW -m tcp -p tcp --dport 5900 -j ACCEPT
```

5. Enter the following command to restart `iptables`:

```
1 # service iptables restart
```

6. Enter the following command to restart `gdm`:

```
1 # telinit 3
2 # telinit 5
```

Note:

Red Hat Linux uses runlevel 5 for graphical startup. If your installation starts up in runlevel 3, change this configuration for the display manager be started and to get access to a graphical console. For more information, see [Check run levels](#).

Set up SLES-based VMs for VNC

Note:

Before setting up your SUSE Linux Enterprise Server VMs for VNC, be sure that you have installed the Citrix VM Tools for Linux. See [Install the Citrix VM Tools for Linux](#) for details.

SLES has support for enabling “Remote Administration” as a configuration option in [YaST](#). You can select to enable Remote Administration at install time, available on the **Network Services** screen of the SLES installer. This feature allows you to connect an external VNC viewer to your guest to allow you to view the graphical console. The method for using the SLES remote administration feature is slightly different than the method provided by XenCenter. However, it is possible to modify the configuration files in your SUSE Linux VM such that it is integrated with the graphical console feature.

Check for a VNC server

Before making configuration changes, verify that you have a VNC server installed. SUSE ships the `tightvnc` server by default. This server is a suitable VNC server, but you can also use the standard RealVNC distribution.

You can check that you have the `tightvnc` software installed by running the command:

```
1 rpm -q tightvnc
```

Enable remote administration

If Remote Administration was not enabled during installation of the SLES software, you can enable it as follows:

1. Open a text console on the VM and run the `YaST` utility:

```
1 yast
```

2. Use the arrow keys to select **Network Services** in the left menu. **Tab** to the right menu and use the arrow keys to select **Remote Administration**. Press **Enter**.
3. In the **Remote Administration** screen, **Tab** to the **Remote Administration Settings** section. Use the arrow keys to select **Allow Remote Administration** and press **Enter** to place an X in the check box.
4. **Tab** to the **Firewall Settings** section. Use the arrow keys to select **Open Port in Firewall** and press **Enter** to place an X in the check box.
5. **Tab** to the **Finish** button and press **Enter**.
6. A message box is displayed, telling you to restart the display manager for your settings to take effect. Press **Enter** to acknowledge the message.
7. The original top-level menu of `YaST` appears. **Tab** to the **Quit** button and press **Enter**.

Modify the `xinetd` configuration

After enabling Remote Administration, modify a configuration file if you want to allow XenCenter to connect. Alternatively, use a third party VNC client.

1. Open the file `/etc/xinetd.d/vnc` in your preferred text editor.
2. The file contains sections like the following:

```
1  service vnc1
2  {
3
4      socket_type = stream
5      protocol   = tcp
6      wait       = no
7      user       = nobody
8      server     = /usr/X11R6/bin/Xvnc
9      server_args = :42 -inetd -once -query localhost -geometry 1024
              x768 -depth 16
10     type      = UNLISTED
11     port      = 5901
12 }
```

3. Edit the `port` line to read

```
1  port = 5900
```

4. Save and close the file.
5. Restart the display manager and `xinetd` service with the following commands:

```
1  /etc/init.d/xinetd restart
2  rcxdm restart
```

SUSE Linux uses runlevel 5 for graphical startup. If your remote desktop does not appear, verify that your VM is configured to start up in runlevel 5. For more information, see [Check Run levels](#).

Firewall settings

By default the firewall configuration does not allow VNC traffic to go through. If you have a firewall between the VM and XenCenter, allow traffic over the port that the VNC connection uses. By default, a VNC server listens for connections from a VNC viewer on TCP port 5900 + `n`, where `n` is the display number (usually zero). So a VNC server setup for Display-0 listens on TCP port 5900, Display-1 is TCP –5901, and so forth. Consult your firewall documentation to ensure that these ports are open.

If you want to use IP connection tracking or limit the initiation of connections to be from one side only, further configure your firewall.

To Open the VNC Port on SLES 11.x VMs Firewall:

1. Open a text console on the VM and run the `YaST` utility:

```
1  yast
```

2. Use the arrow keys to select **Security and Users** in the left menu. **Tab** to the right menu and use the arrow keys to select **Firewall**. Press **Enter**.

3. In the **Firewall** screen, use the arrow keys to select **Custom Rules** in the left menu and then press **Enter**.
4. **Tab** to the **Add** button in the **Custom Allowed Rules** section and then press **Enter**.
5. In the **Source Network** field, enter `0/0`. **Tab** to the **Destination Port** field and enter `5900`.
6. **Tab** to the **Add** button and then press **Enter**.
7. **Tab** to the **Next** button and press **Enter**.
8. In the **Summary** screen **Tab** to the **Finish** button and press **Enter**.
9. On the top-level **YaST** screen **Tab** to the **Quit** button and press **Enter**.
10. Restart the display manager and `xinetd` service with the following commands:

```
1 /etc/init.d/xinetd restart
2 rcxdm restart
```

Alternatively, you can disable the firewall until the next reboot by running the **rcSuSEfirewall2 stop** command, or permanently by using **YaST**. This configuration can expose extra services to the outside world and reduce the overall security of your VM.

VNC screen resolution

After connecting to a Virtual Machine with the Graphical Console, the screen resolution sometimes does not match. For example, the VM display is too large to fit comfortably in the Graphical Console pane. Control this behavior by setting the VNC server `geometry` parameter as follows:

1. Open the `/etc/xinetd.d/vnc` file with your preferred text editor and find the `service_vnc1` section (corresponding to `displayID 1`).
2. Edit the `geometry` argument in the `server-args` line to the desired display resolution. For example,

```
1 server_args = :42 -inetd -once -query localhost -geometry 800x600
   -depth 16
```

The value of the `geometry` parameter can be any valid screen width and height.

3. Save and close the file.
4. Restart the VNC server:

```
1 /etc/init.d/xinetd restart
2 rcxdm restart
```

Check run levels

Red Hat and SUSE Linux VMs use runlevel 5 for graphical startup. This section describes how to verify that your VM starts up in runlevel 5 and how to change this setting.

1. Check `/etc/inittab` to see what the default runlevel is set to. Look for the line that reads:

```
1 id:n:initdefault:
```

If *n* is not 5, edit the file to make it so.

2. You can run the command `telinit q ; telinit 5` after this change to avoid having to reboot to switch run levels.

Troubleshoot VM problems

May 13, 2024

Citrix provides two forms of support:

- Free, self-help support on the [Citrix website](#)
- Paid-for Support Services, which you can buy from the Support Site.

With Citrix Technical Support, you can open a Support Case online or contact the support center by phone if you experience technical difficulties.

The [Citrix Support](#) site hosts several resources that might be helpful to you if you experience unusual behavior, crashes, or other problems. Resources include: Support Forums, Knowledge Base articles, and product documentation.

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 Citrix Hypervisor Solution Provider track and resolve the issue.

Important:

Follow the troubleshooting information in this section only under the guidance of your Citrix Hypervisor 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 Citrix Hypervisor server, 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

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 the **Update & Security** button
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. In Citrix Hypervisor 8.2, only the following operating systems support Secure Boot: Windows 10 (64-bit), Windows Server 2016 (64-bit), Windows Server 2019 (64-bit), Windows Server 2022 (64-bit).

Check that your Citrix Hypervisor server is booted in UEFI mode. You can only create UEFI Secure Boot VMs on a Citrix Hypervisor server that has the Secure Boot certificates present. Secure Boot certificates are only present on servers booted in UEFI mode or servers in the same pool as a server booted in UEFI mode. For more information, see [Network Boot](#).

Check that the UEFI-booted Citrix Hypervisor server is included in the [Hardware Compatibility List](#). Older servers might not include the Secure Boot certificates when booted in UEFI mode.

How do I know if a Citrix Hypervisor server has the Secure Boot certificates?

If your Citrix Hypervisor server is booted in UEFI mode, the Secure Boot certificates are available on the server. Citrix Hypervisor servers share their certificates with other servers in the same resource pool. If you have a UEFI booted server in your resource pool, all servers in that pool have the Secure Boot certificates available.

Run the following command on a Citrix Hypervisor server:

```
1 xe pool-param-get param-name=uefi-certificates uuid=<pool-uuid> | wc -c
```

If it returns a value that is greater than zero, the Secure Boot certificates are present.

To check that the certificates are valid, run the following command on your Citrix Hypervisor server:

```
1 xe pool-param-get uuid=$(xe pool-list --minimal) param-name=uefi-  
  certificates|base64 -d|tar tv  
2 -rw-r--r-- root/root      1600 2019-11-11 17:09 KEK.auth  
3 -rw-r--r-- root/root      3212 2019-11-11 17:09 db.auth
```

If the Secure Boot certificates are absent, run the following command on your Citrix Hypervisor server:

```
1 ls /sys/firmware/efi/efivars | grep KEK
```

If this command returns empty, Secure Boot VMs cannot be created on that server because the required certificates are missing from the UEFI firmware.

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 and start the VM in setup mode to remove the unsigned drivers.

Important:

Before doing this, back up your VM by taking a snapshot.

To change a UEFI Secure Boot VM into a UEFI boot VM, run the following command on the Citrix Hypervisor server that hosts the VM:

```
1 varstore-sb-state <VM_UUID> setup
```

After you have fixed your VM, run the following command to re-enable Secure Boot:

```
1 varstore-sb-state <VM_UUID> user
```

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.

To disable Secure Boot, run the following command on the Citrix Hypervisor server that hosts the VM:

```
1 varstore-sb-state <VM_UUID> setup
```

After you have debugged the issue, you can run the following command to re-enable Secure Boot:

```
1 varstore-sb-state <VM_UUID> user
```

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 can do one of the following things:

- Switch your VM to UEFI boot mode by running the following command:

```
1 xe vm-param-set uuid=<UUID> platform:secureboot=false
```

Reboot the VM.

After you have debugged the issue, you can run the following command to re-enable Secure Boot:

```
1 xe vm-param-set uuid=<UUID> platform:secureboot=auto
```

Reboot the VM.

- Disable Secure Boot by running the following command on the Citrix Hypervisor server that hosts the VM:

```
1 varstore-sb-state <VM_UUID> setup
```

After you have debugged the issue, you can run the following command to re-enable Secure Boot:

```
1 varstore-sb-state <VM_UUID> user
```

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. This information displays correctly after the XenServer VM Tools for Windows have been installed in the VM.

Why are my emulated devices showing as different types than expected on a UEFI Windows VM?

A UEFI Secure Boot VMs use NVME and E1000 for emulated devices. However, when the VM first starts the emulated devices show as different types. This information displays correctly after the XenServer VM Tools for Windows have been installed in the VM.

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 Citrix Hypervisor.

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 Citrix Hypervisor server 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 Citrix Hypervisor 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 Citrix Hypervisor 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 server in a table. The information is periodically refreshed.

Output columns

The xentop utility displays the following columns in the console:

- **NAME** - The name of the domain. “Domain-0” is the Citrix Hypervisor 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

July 15, 2024

High availability is a set of automatic features designed to plan for, and safely recover from issues which take down Citrix Hypervisor servers or make them unreachable. For example, during physically disrupted networking or host hardware failures.

Overview

High availability ensures that if a host becomes unreachable or unstable, the VMs running on that host are safely restarted on another host automatically. This removes the need for the VMs to be manually restarted, resulting in minimal VM downtime.

When the pool master becomes unreachable or unstable, high availability can also recover administrative control of a pool. High availability ensures that administrative control is restored automatically without any manual intervention.

Optionally, high availability can also automate the process of restarting VMs on hosts which are known to be in a good state without manual intervention. These VMs can be scheduled for restart in groups to allow time to start services. It allows infrastructure VMs to be started before their dependent VMs (for example, a DHCP server before its dependent SQL server).

Warnings:

Use high availability along with multipathed storage and bonded networking. Configure multipathed storage and bonded networking before attempting to set up high availability. Customers who do not set up multipathed storage and bonded networking can see unexpected host reboot behavior (Self-Fencing) when there is an infrastructure instability.

All graphics solutions (NVIDIA vGPU, Intel GVT-d, Intel GVT-G, AMD MxGPU, 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.

Overcommitting

A pool is overcommitted when the VMs that are currently running cannot be restarted elsewhere following a user-defined number of host failures. Overcommitting can happen if there is not enough free memory across the pool to run those VMs following a failure. However, there are also more subtle changes which can make high availability unsustainable: changes to Virtual Block Devices (VBDs) and networks can affect which VMs can be restarted on which hosts. Citrix Hypervisor cannot check all potential actions and determine if they cause violation of high availability demands. However, an asynchronous notification is sent if high availability becomes unsustainable.

Citrix Hypervisor dynamically maintains a failover plan which details what to do when a set of hosts in a pool fail at any given time. An important concept to understand is the `host failures to tolerate` value, which is defined as part of the high availability configuration. The value of `host failures to tolerate` determines the number of host failures that are allowed while still being able to restart all protected VMs. For example, consider a resource pool that consists of 64 hosts and `host failures to tolerate` is set to 3. In this case, the pool calculates a failover plan that allows any three hosts to fail and then restarts the VMs on other hosts. If a plan cannot be found, then the pool is considered to be overcommitted. The plan is dynamically recalculated based on VM lifecycle operations and movement. If changes (for example, the addition of new VMs to the pool) cause the pool to become overcommitted, alerts are sent either via XenCenter or email.

Overcommitment warning

If any attempts to start or resume a VM would cause the pool to become overcommitted, a warning alert is displayed in XenCenter. You can then choose to cancel the operation or proceed anyway. Proceeding causes the pool to become overcommitted and a message is sent to any configured email addresses. This is also available as a message instance through the management API. The amount of memory used by VMs of different priorities is displayed at the pool and host levels.

Host fencing

Sometimes a server can fail due to the loss of network connectivity or when a problem with the control stack is encountered. In such cases, the Citrix Hypervisor server self-fences to ensure that the VMs are not running on two servers simultaneously. When a fence action is taken, the server restarts immediately and abruptly, causing all VMs running on it to be stopped. The other servers detect that the VMs are no longer running and then the VMs are restarted according to the restart priorities assigned to them. The fenced server enters a reboot sequence, and when it has restarted it tries to rejoin the resource pool.

Note:

Hosts in clustered pools can also self-fence when they cannot communicate with more than half of the other hosts in the resource pool. For more information, see [Clustered pools](#).

Configuration requirements

To use the high availability feature, you need:

- Citrix Hypervisor pool (this feature provides high availability at the server level within a single resource pool).

Note:

We recommend that you enable high availability only in pools that contain at least 3 Citrix Hypervisor servers. For more information, see [CTX129721 - High Availability Behavior When the Heartbeat is Lost in a Pool](#).

- Shared storage, including at least one iSCSI, NFS, or Fibre Channel LUN of size 356 MB or greater - the *heartbeat SR*. The high availability mechanism creates two volumes on the heartbeat SR:
 - 4 MB heartbeat volume: Used to provide a heartbeat.
 - 256 MB metadata volume: To store pool master metadata to be used if there is a master failover.

Note:

Previously, we recommended using a dedicated NFS or iSCSI storage repository as your high availability heartbeat disk. However, this is only of benefit if the storage repository is not sharing resources on the underlying storage appliance, otherwise it is simply increasing complexity and resource usage in the control domain (dom0) of the host.

If your pool is a clustered pool, your heartbeat SR must be a GFS2 SR.

Storage attached using either SMB or iSCSI when authenticated using CHAP cannot be used

as the heartbeat SR.

- Static IP addresses for all hosts.

Warning:

If the IP address of a server changes while high availability is enabled, high availability assumes that the host's network has failed. The change in IP address can fence the host and leave it in an unbootable state. To remedy this situation, disable high availability using the `host-emergency-ha-disable` command, reset the pool master using `pool-emergency-reset-master`, and then re-enable high availability.

- For maximum reliability, we recommend that you use a dedicated bonded interface as the high availability management network.

For a VM to be protected by high availability, it must be agile. This means the VM:

- Must have its virtual disks on shared storage. You can use any type of shared storage. iSCSI, NFS, or Fibre Channel LUN is only required for the storage heartbeat and can be used for virtual disk storage.
- Can use live migration.
- Does not have a connection to a local DVD drive configured.
- Has its virtual network interfaces on pool-wide networks.

Note:

When high availability is enabled, we strongly recommend using a bonded management interface on the servers in the pool and multipathed storage for the heartbeat SR.

If you create VLANs and bonded interfaces from the CLI, then they might not be plugged in and active despite being created. In this situation, a VM can appear to be not agile and it is not protected by high availability. You can use the CLI `pif-plug` command to bring up the VLAN and bond PIFs so that the VM can become agile. You can also determine precisely why a VM is not agile by using the `xe diagnostic-vm-status` CLI command. This command analyzes its placement constraints, and you can take remedial action if necessary.

Restart configuration settings

Virtual machines can be considered protected, best-effort, or unprotected by high availability. The value of `ha-restart-priority` defines whether a VM is treated as protected, best-effort, or unprotected. The restart behavior for VMs in each of these categories is different.

Protected

High availability guarantees to restart a protected VM that goes offline or whose host goes offline, provided the pool isn't overcommitted and the VM is agile.

If a protected VM cannot be restarted when a server fails, high availability attempts to start the VM when there is extra capacity in a pool. Attempts to start the VM when there is extra capacity might now succeed.

`ha-restart-priority` Value: `restart`

Best-effort

If the host of a best-effort VM goes offline, high availability attempts to restart the best-effort 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.

`ha-restart-priority` Value: `best-effort`

Unprotected

If an unprotected VM or the host it runs on is stopped, high availability does not attempt to restart the VM.

`ha-restart-priority` Value: Value is an empty string

Note:

High availability never stops or migrates a running VM to free resources for a protected or best-effort VM to be restarted.

If the pool experiences server failures and the number of tolerable failures drops to zero, the protected VMs are not guaranteed to restart. In such cases, a system alert is generated. If another failure occurs, all VMs that have a restart priority set behave according to the best-effort behavior.

Start order

The start order is the order in which Citrix Hypervisor high availability attempts to restart protected VMs when a failure occurs. The values of the `order` property for each of the protected VMs determines the start order.

The `order` property of a VM is used by high availability and also by other features that start and shut down VMs. Any VM can have the `order` property set, not just the VMs marked as protected for high availability. However, high availability uses the `order` property for protected VMs only.

The value of the `order` property is an integer. The default value is 0, which is the highest priority. Protected VMs with an `order` value of 0 are restarted first by high availability. The higher the value of the `order` property, the later in the sequence the VM is restarted.

You can set the value of the `order` property of a VM by using the command-line interface:

```
1 xe vm-param-set uuid=<vm uuid> order=<int>
```

Or in XenCenter, in the **Start Options** panel for a VM, set **Start order** to the required value.

Enable high availability on your Citrix Hypervisor pool

You can enable high availability on a pool by using either XenCenter or the command-line interface (CLI). In either case, you specify a set of priorities that determine which VMs are given the highest restart priority when a pool is overcommitted.

Warnings:

- When you enable high availability, some operations that compromise the plan for restarting VMs (such as removing a server from a pool) might be disabled. You can temporarily disable high availability to perform such operations, or alternatively, make VMs protected by high availability unprotected.
- If high availability is enabled, you cannot enable clustering on your pool. Temporarily disable high availability to enable clustering. You can enable high availability on your clustered pool. Some high availability behavior, such as self-fencing, is different for clustered pools. For more information, see [Clustered pools](#).

Enable high availability by using the CLI

1. Verify that you have a compatible Storage Repository (SR) attached to your pool. iSCSI, NFS, or Fibre Channel SRs are compatible. For information about how to configure such a storage repository using the CLI, see [Manage storage repositories](#).
2. For each VM you want to protect, set a restart priority and start order. You can set the restart priority as follows:

```
1 xe vm-param-set uuid=<vm uuid> ha-restart-priority=restart order=1
```

3. Enable high availability on the pool, and optionally, specify a timeout:

```
1 xe pool-ha-enable heartbeat-sr-uuids=<sr uuid> ha-config:timeout=<
  timeout in seconds>
```

Alternatively, you can set a default timeout for your pool. For more information about how to set a timeout, see [Configure high availability timeout](#).

4. Run the command `pool-ha-compute-max-host-failures-to-tolerate`. This command returns the maximum number of hosts that can fail before there are insufficient resources to run all the protected VMs in the pool.

```
1 xe pool-ha-compute-max-host-failures-to-tolerate
```

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 for `ha-host-failures-to-tolerate`.

5. Specify the `ha-host-failures-to-tolerate` parameter. The value must be less than or equal to the computed value:

```
1 xe pool-param-set ha-host-failures-to-tolerate=2 uuid=<pool uuid>
```

Configure high availability timeout

High availability timeout is the period during which networking or storage is not accessible by the hosts in your pool. If any Citrix Hypervisor server 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 using the following command.

Set 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.

Alternatively, you can set timeout when you enable high availability:

```
1 xe pool-ha-enable heartbeat-sr-uuids=<sr uuid> ha-config:timeout=<
  timeout in seconds>
```

Note that if you set timeout when enabling high availability, it only applies to that specific enablement. Therefore, if you disable and then reenabling high availability later, the high availability feature reverts to using the default timeout.

Remove high availability protection from a VM by using the CLI

To disable high availability features for a VM, use the `xe vm-param-set` command to set the `ha-restart-priority` parameter to be an empty string. Setting the `ha-restart-priority` parameter does not clear the start order settings. You can enable high availability for a VM again by setting the `ha-restart-priority` parameter to `restart` or `best-effort` as appropriate.

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 Citrix Hypervisor installation, you might have to disable high availability using the `host-emergency-ha-disable` command:

```
1 xe host-emergency-ha-disable --force
```

If the host was the pool master, 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 master, you might have to take one of the following actions:

- Force the host to reboot as a pool master (`xe pool-emergency-transition-to-master`)

```
1 xe pool-emergency-transition-to-master uuid=<host uuid>
```

- Tell the host where the new master is (`xe pool-emergency-reset-master`):

```
1 xe pool-emergency-reset-master master-address=<new master  
hostname>
```

When all hosts have successfully restarted, re-enable high availability:

```
1 xe pool-ha-enable heartbeat-sr-uuid=<sr uuid>
```

Shutting 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. XenCenter offers you a dialog box to automate disabling the protection when you select the **Shutdown** 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 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 7, 2022

The Citrix Hypervisor 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 server 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 Citrix Hypervisor DR

Citrix Hypervisor 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.

Citrix Hypervisor 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 server 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 Citrix Hypervisor 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 Citrix Hypervisor environment must be at the same release and patch level.
- The servers and pools at the secondary site must have the same license edition as those at the primary site. These Citrix Hypervisor licenses are in addition to those assigned to servers at the primary site.

If you have a Citrix Virtual Apps and Desktops entitlement or a Citrix DaaS entitlement, you can use the same entitlement for both your primary and secondary 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

January 10, 2023

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 prechecks.

4. The wizard performs several prechecks 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 prechecks 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 prechecks.

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 prechecks 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 prechecks.
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 prechecks 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 7, 2022

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 you to 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 servers 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

May 28, 2024

Whenever possible, leave the installed state of Citrix Hypervisor servers unaltered. That is, do not install any additional packages or start additional services on Citrix Hypervisor servers and treat them as appliances. The best way to restore, then, is to reinstall Citrix Hypervisor server software from the installation media. If you have multiple Citrix Hypervisor servers, 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 one of our certified partners. For more information, see [Citrix Ready Marketplace](#).

These solutions are certified by their vendors and not by Cloud Software Group. Support for these backup solutions is provided by the solution vendor.

Citrix Hypervisor Premium Edition customers can take advantage of the faster changed block only backup. For more information, see the Citrix blog about [Changed Block Tracking backup APIs](#).

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

Citrix Hypervisor servers 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 Citrix Hypervisor server 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 Citrix Hypervisor server has died completely, then you must first do a fresh install, and then run the `pool-restore-database` command against the freshly installed Citrix Hypervisor server.

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 **Halting** 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:

Citrix Hypervisor preserves UUIDs of the hosts restored using this method. If you restore to a different physical machine while the original Citrix Hypervisor server is still running, duplicate UUIDs may be present. As a result, XenCenter refuses to connect to the second Citrix Hypervisor server. 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 master host 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 master 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 master, 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 master host, and archive this file. To restore this backup later on a brand new set of hosts:

1. Install a fresh set of Citrix Hypervisor servers 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 master.
3. Run the `xe host-forget` command on the new master 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 Citrix Hypervisor servers

This section describes the Citrix Hypervisor server 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 Citrix Hypervisor 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 Citrix Hypervisor cleanly from the Citrix Hypervisor 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 Citrix Hypervisor server:

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 Citrix Hypervisor server:

1. If you want to restore your Citrix Hypervisor server from a specific backup, run the following command while the Citrix Hypervisor server 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 Citrix Hypervisor server 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 Citrix Hypervisor server using the Citrix Hypervisor installation CD and select the **Restore from backup** option.

After the restore from backup is completed, reboot the Citrix Hypervisor server 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 Citrix Hypervisor server:

If your Citrix Hypervisor server has crashed and is not reachable, use the Citrix Hypervisor 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 Citrix Hypervisor servers as describes in this section.

Back up VMs

We recommend that you use a backup solution offered by one of our certified partners. For more information, see [Citrix Ready Marketplace](#).

These solutions are certified by their vendors and not by Cloud Software Group. Support for these backup solutions is provided by the solution vendor.

Citrix Hypervisor Premium Edition customers can take advantage of the faster changed block only backup. For more information, see the Citrix blog about [Changed Block Tracking backup APIs](#).

VM snapshots

March 13, 2024

Citrix Hypervisor provides a convenient mechanism that can take a snapshot of a VM storage and meta-data 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 Citrix Hypervisor, 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 by using 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 Citrix Hypervisor 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 Citrix Hypervisor pool

Run the `snapshot-list` command:

```
1 xe snapshot-list
```

This command lists all of the snapshots in the Citrix Hypervisor 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 ( R0): 116dd310-a0ef-a830-37c8-df41521ff72d  
3 name-label ( RW): Windows Server 2016 (1)  
4 power-state ( R0): halted  
5  
6 uuid ( R0): dff45c56-426a-4450-a094-d3bba0a2ba3f  
7 name-label ( RW): Control domain on host  
8 power-state ( R0): 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 re-

store 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's 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 Citrix Hypervisor 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 Citrix Hypervisor server.

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 server.

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

March 1, 2023

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, master nodes detect the failures of members by receiving regular heartbeat messages. If no heartbeat has been received for 600 seconds, the master 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 master marks the member as alive again.
- Shut down the host and instruct the master 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 Citrix Hypervisor servers.

It is important to ensure that the Citrix Hypervisor server 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 Citrix Hypervisor 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 Citrix Hypervisor server fails, there might be VMs still registered in the *running* state. If you are sure that the member Citrix Hypervisor server 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 Citrix Hypervisor server, 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 Citrix Hypervisor servers once a host has failed. To do so, run the following script on the pool master 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 master at the time of the crash. (The SR master is the pool master or a Citrix Hypervisor server 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 Citrix Hypervisor server before running the `resetvdis.py` script, then you receive the following error message: `VDI <UUID> already attached RW`.

Master failures

Every member of a resource pool contains all the information necessary to take over the role of master if necessary. When a master node fails, the following sequence of events occurs:

1. If HA is enabled, another master is elected automatically.
2. If HA is not enabled, each member waits for the master to return.

If the master comes back up at this point, it re-establishes communication with its members, and operation returns to normal.

If the master is dead, choose one of the members and run the command `xe pool-emergency-transition-to-master` on it. Once it has become the master, run the command `xe pool-recover-slaves` and the members now point to the new master.

If you repair or replace the server that was the original master, you can simply bring it up, install the Citrix Hypervisor software, and add it to the pool. Since the Citrix Hypervisor servers in the pool are enforced to be homogeneous, there is no real need to make the replaced server the master.

When a member Citrix Hypervisor server is transitioned to being a master, 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 master, restore the pool database from your backup using the `xe pool-restore-database` command (see `pool-restore-database`).

3. Connect to the master host 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 master 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.

Troubleshooting

July 15, 2024

Support

Citrix provides two forms of support: free, self-help support on the [Citrix Support](#) website and paid-for Support Services, which you can purchase from the Support site. With Citrix Technical Support, you can open a Support Case online or contact the support center by phone if you experience technical difficulties.

The [Citrix Knowledge Center](#) hosts several resources that might be helpful to you in the event of odd behavior, crashes, or other problems. Resources include: Forums, Knowledge Base articles, White Papers, product documentation, hotfixes, and other updates.

If you experience technical difficulties with the Citrix Hypervisor server, 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 your Solution Provider track and resolve the issue.

For information about troubleshooting Citrix Hypervisor installation issues, see [Troubleshoot the installation](#). For information about troubleshooting virtual machine issues, see [Troubleshoot VM problems](#).

Important:

We recommend that you follow the troubleshooting information in this section solely under the guidance of your Solution Provider or the Support team.

In some support cases, serial console access is required for debug purposes. Therefore, when setting up your Citrix Hypervisor configuration, it is recommended that serial console access is configured. For hosts that do not have physical serial port (such as a Blade server) or where suitable physical infrastructure is not available, investigate whether an embedded management device, such as Dell DRAC can be configured.

For information on setting up serial console access, see [CTX228930 - How to configure console access on XenServer or Citrix Hypervisor](#).

Citrix Hypervisor server logs

Citrix Hypervisor server status reports

XenCenter can be used to gather Citrix Hypervisor server information.

Click **Server Status Report** in the **Tools** menu to open the **Server Status Report** task. You can select from a list of different types of information (various logs, crash dumps, and so on). The information is compiled and downloaded to the machine that XenCenter is running on. For more information, see the [XenCenter documentation](#).

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 command `xenserver-status-report -u` in the Citrix Hypervisor server console.

Important:

Citrix Hypervisor server logs may contain sensitive information.

Sending host log messages to a central server

Rather than have logs written to the control domain filesystem, you can configure your Citrix Hypervisor server 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 hostname or IP address of the remote server where you want the logs to be written:

```
1 xe host-param-set uuid=host_uuid logging:syslog_destination=hostname
```

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 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. Depending on your version of Windows, the path can be one of the following:

- `%userprofile%\AppData\Citrix\XenCenter\logs\XenCenter.log`
- `%userprofile%\AppData\Citrix\Roaming\XenCenter\logs\XenCenter.log`

To locate the XenCenter log files - for example, when you want to open or email the log file - click **View Application Log Files** in the XenCenter **Help** menu.

Troubleshooting connections between XenCenter and the Citrix Hypervisor server

If you have trouble connecting to the Citrix Hypervisor server with XenCenter, check the following:

- Is your XenCenter an older version than the Citrix Hypervisor server you are attempting to connect to?

The XenCenter application is backward-compatible and can communicate properly with older Citrix Hypervisor servers, but an older XenCenter cannot communicate properly with newer Citrix Hypervisor servers.

To correct this issue, install the XenCenter version that is the same, or newer, than the Citrix Hypervisor server version.

- Is your license current?

You can see the expiration date for your license access code in the Citrix Hypervisor server **General** tab under the **License Details** section in XenCenter.

For more information on licensing a host, see [Licensing](#).

- The Citrix Hypervisor server 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 Citrix Hypervisor server and the machine running the client software, ensure that it allows traffic from these ports.

Other troubleshooting information

The following articles provide troubleshooting information about specific areas of the product:

- [Install troubleshooting](#)
- [VM troubleshooting](#)
- [Networking troubleshooting](#)

Workload Balancing

June 11, 2024

Notes:

- Workload Balancing is available for Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, see [Licensing](#). To upgrade, or to get a Citrix Hypervisor license, visit the [Citrix website](#).
- Workload Balancing 8.2.2 and later are compatible with Citrix Hypervisor 8.2 Cumulative Update 1.

Workload Balancing is a Citrix Hypervisor Premium Edition component, packaged as a virtual appliance, that provides the following features:

- Create reports about virtual machine (VM) performance in your Citrix Hypervisor 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 Citrix Hypervisor 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 servers 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

July 22, 2024

The latest version of the Workload Balancing virtual appliance is version 8.3.0. You can download this version of the Workload Balancing virtual appliance from the [Citrix Hypervisor Product Software downloads page](#).

What's new in 8.3.0

Released Feb 23, 2023

This update includes the following improvements:

- Via XenAPI, you can now set the alert level for Workload Balancing alerts in XenCenter.

This update includes changes to the WLB database. Ensure that you use the provided migration script when you update your WLB to this version. For more information about using the migration script, see [Migrate data from an existing virtual appliance](#).

Fixed issues in 8.3.0

This update fixes the following issues:

- During the Workload Balancing maintenance window, Workload Balancing is unable to provide placement recommendations. When this situation occurs, you see the error: “4010 Pool discovery has not been completed. Using original algorithm.”The Workload Balancing maintenance window is less than 20 minutes long and by default is scheduled at midnight.
- For a Workload Balancing virtual appliance version 8.2.2 and later that doesn't use LVM, you cannot extend the available disk space.
- Due to an unresponsive API call, Workload Balancing is sometimes blocked during pool discovery.
- In XenCenter, the date range and some timestamps shown on the Workload Balancing Pool Audit Report are incorrect.
- In XenCenter, some strings are not displaying correctly for Workload Reports.
- If the Workload Balancing virtual appliance is running for a long time, it is shut down by the operating system for consuming a lot of memory.
- The database fails to auto-restart after the Workload Balancing virtual appliance experiences an abnormal shutdown.

Known issues in 8.3.0

This update contains the following known issues:

- If you do not use the default user name ([postgres](#)) for the Workload Balancing PostgreSQL database, you cannot connect your pools to the Workload Balancing virtual appliance.

To work around this issue, use the default user name [postgres](#) for the Workload Balancing database.

- After configuring Workload Balancing on a pool, if your host CPU load is low, you might encounter an [Internal error](#) when attempting to put your host into maintenance mode.

To work around this issue, pause the Workload Balancing virtual appliance or disconnect it from your pools.

Earlier releases

This section lists features in previous releases along with their fixed issues. These earlier releases are superseded by the latest version of the Workload Balancing virtual appliance. Update to the latest version of the Workload Balancing virtual appliance when it is available.

Workload Balancing 8.2.2

Released Sep 30, 2021

This update includes changes to the WLB database. Ensure that you use the provided migration script when you update your WLB to this version. For more information about using the migration script, see [Migrate data from an existing virtual appliance](#).

Fixed issues This update includes fixes for the following issues:

- The Workload Balancing database can grow very fast and fill the disk.
- A race condition can sometimes cause records to be duplicated in the WLB database. When this occurs, the user might see the error: “WLB received an unknown exception”.

Workload Balancing 8.2.1

Released Sep 15, 2020

This update includes the following improvements:

- The migration script now enables you to migrate your Workload Balancing database from the Workload Balancing virtual appliance 8.0.0 (which was provided with Citrix Hypervisor 8.0 and 8.1) to the Workload Balancing virtual appliance 8.2.1 provided with Citrix Hypervisor 8.2.

For more information about using the migration script, see [Migrate data from an existing virtual appliance](#).

Fixed issues This update includes fixes for the following issues:

- When multiple VMs start at the same time, Workload Balancing recommends balancing the VMs placement on all servers in the pool evenly. However, sometimes Workload Balancing might recommend to put many VMs on the same Citrix Hypervisor server. This issue occurs when Workload Balancing gets late feedback from XAPI about VM placement.

Get started with Workload Balancing

July 1, 2024

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.](#)
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 Citrix Hypervisor server. Before importing it, review the prerequisite information and considerations.

Prerequisites

- Workload Balancing 8.2.2 and later are compatible with Citrix Hypervisor 8.2 Cumulative Update 1. 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 migrate 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 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. For more information, see [Change the Workload Balancing virtual appliance configuration](#).

Pool requirements

To balance a pool with Workload Balancing, the pool must meet the following requirements:

- All servers are licensed with a Premium Edition license or a Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement
- All servers 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.

A single Workload Balancing virtual appliance can manage multiple pools up to a maximum of 100 pools, depending on the virtual appliance's resources (vCPU, memory, disk size). Across these pools, the virtual appliance can manage up to 1000 VMs. However, if a 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.

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 Citrix Hypervisor.

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 Citrix Hypervisor server 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 Citrix Hypervisor pool, you must specify an existing account:

- *Citrix Hypervisor account*

This account is used by the Workload Balancing virtual appliance to connect to the Citrix Hypervisor pool and read the RRDs. Ensure that this user account has the permissions to read the Citrix Hypervisor pool, server, 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.
- **Citrix Hypervisor 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 Citrix Hypervisor 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 2,048 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 Citrix download page <http://www.citrix.com/downloads>. 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 Citrix Hypervisor:

1. Open XenCenter.
2. Right-click on the pool (or server) 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. Citrix recommends selecting a strong password.

Note:

When you enter the password, the console does not display placeholders, such as asterisks, for the characters.

4. Enter the computer name you want to assign to the Workload Balancing 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 `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, subnet mask, and 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 user name for the Workload Balancing database, or press **Enter** to use the default user name (`postgres`) of the database account.

Note:

We recommend to use the default user name `postgres` due to a known issue where you cannot connect your pools to the Workload Balancing virtual appliance if you do not use the default user name for the Workload Balancing PostgreSQL database.

You are creating an account for the Workload Balancing database. The Workload Balancing services use this account to read/write to the Workload Balancing database. Note the user name and password. You might need them if you ever want to administer to the Workload Balancing PostgreSQL database directly (for example, if you wanted to export data).

8. Enter a password for the Workload Balancing database. 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 Citrix Hypervisor 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:

A single Workload Balancing virtual appliance can manage multiple pools up to a maximum of 100 pools, depending on the virtual appliance's resources (vCPU, memory, disk size). Across

these pools, the virtual appliance can manage up to 1000 VMs. However, if a 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.

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, Citrix Hypervisor 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. Citrix Hypervisor 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 master) you want Workload Balancing to monitor.

This account is used by the Workload Balancing virtual appliance to connect to the Citrix Hypervisor pool. This account is created on the Citrix Hypervisor pool master 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

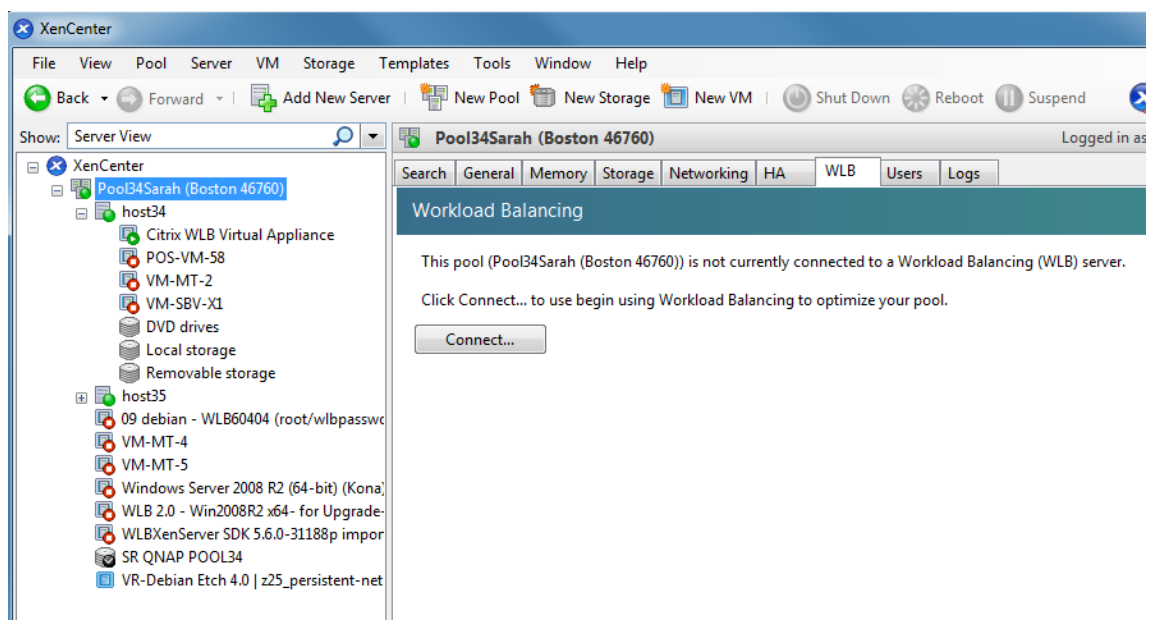
If you want to upload a different (trusted) certificate or configure certificate verification, note the following before connecting your pool to Workload Balancing:

- If you want Citrix Hypervisor 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.

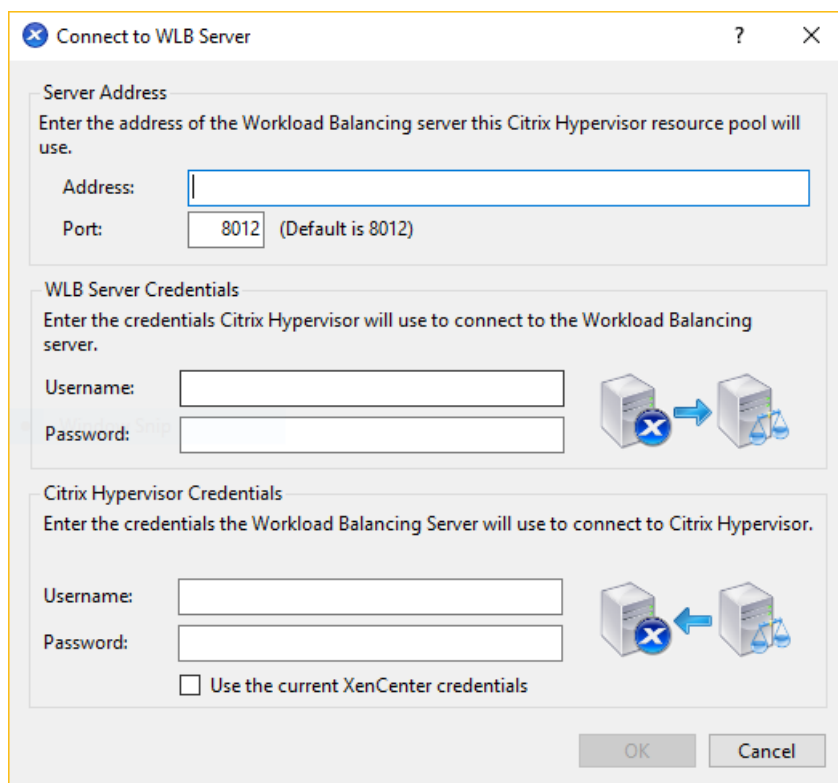
For more information, see [Certificates](#).

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

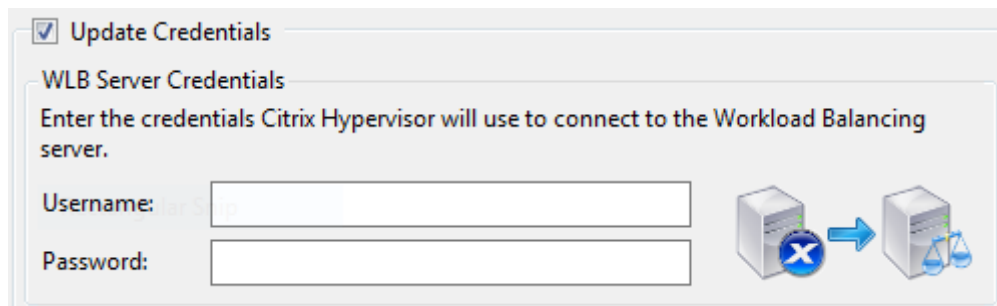
OK Cancel

3. In the **Server Address** section, enter the following:

- In the **Address** box, type the IP address or FQDN of the Workload Balancing virtual appliance. For example, `WLB-appliance-computername.yourdomain.net`.
- (Optional) If you changed the Workload Balancing port during Workload Balancing Configuration, enter the port number in the **Port** box. Citrix Hypervisor uses this port to communicate with Workload Balancing.

By default, Citrix Hypervisor 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**

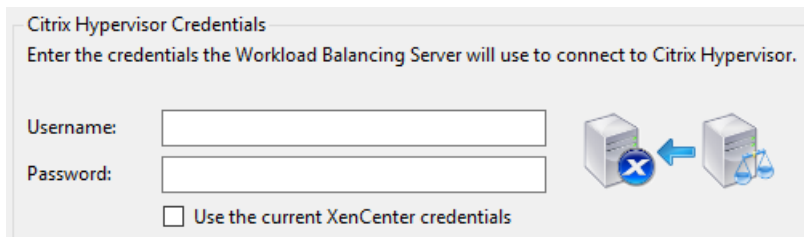
WLB Server Credentials
Enter the credentials Citrix Hypervisor will use to connect to the Workload Balancing server.

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 **Citrix Hypervisor Credentials** section, enter the user name and password for the pool you are configuring. Workload Balancing uses these credentials to connect to the servers in the pool.

A screenshot of the 'Citrix Hypervisor Credentials' dialog box. The title bar says 'Citrix Hypervisor Credentials'. Below the title bar, it says '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 showing 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 Citrix Hypervisor, 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 Citrix Hypervisor 8.2, you can use the migrate script to migrate your existing data when you upgrade to the latest version (Workload Balancing 8.2.1 or later).

The version of Workload Balancing currently provided with Citrix Hypervisor 8.2 is 8.3.0. However, Workload Balancing 8.2.0, 8.2.1, and 8.2.2 were previously available with Citrix Hypervisor 8.2. You can also use this migration script to migrate from Workload Balancing 8.2.1 or 8.2.2 to Workload Balancing 8.3.0.

To use the migrate 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 8.2.1 virtual appliance, run:

```
1 /opt/vpx/wlb/migrate_db.sh 8.2.1 <IP of existing Workload  
Balancing appliance>
```

- For the Workload Balancing 8.2.2 virtual appliance, run:

```
1 /opt/vpx/wlb/migrate_db.sh 8.2.2 <IP of existing Workload  
Balancing appliance>
```

The command prompts you for password information when required.

3. Connect the Citrix Hypervisor 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 Citrix Hypervisor pool.

Workload Balancing basic tasks

July 18, 2024

When you first begin using Workload Balancing, there are some basic tasks you use Workload Balancing for regularly:

- [Determining the best server 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 Citrix Hypervisor 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.

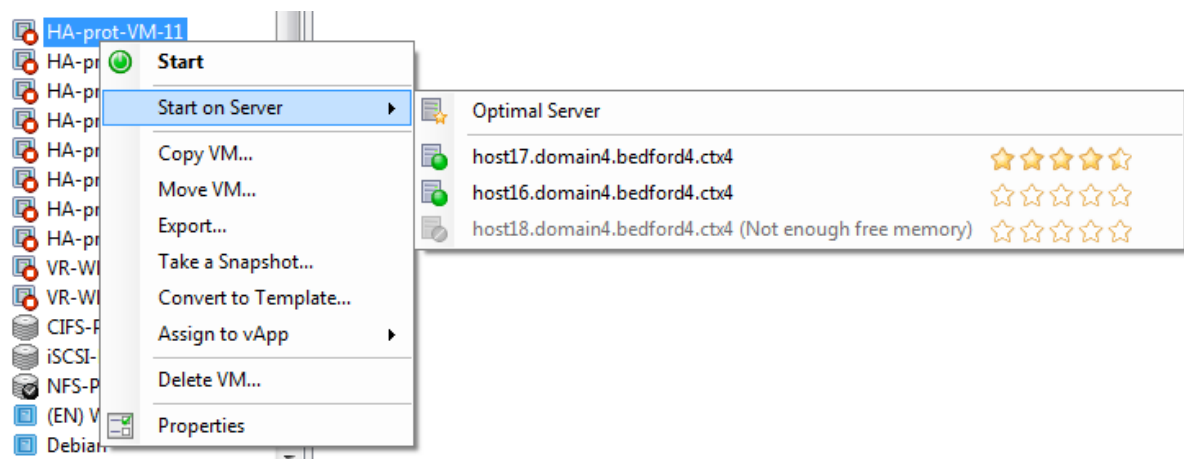
For more information about these more complex features, see [Administer Workload Balancing](#).

Notes:

- Workload Balancing is available for Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, see [Licensing](#). To upgrade, or to get a Citrix Hypervisor license, visit the [Citrix website](#).
- Workload Balancing 8.2.2 and later are compatible with Citrix Hypervisor 8.2 Cumulative Update 1.
- The base OS of the Workload Balancing virtual appliance was originally derived from a general-purpose Linux distribution. The Workload Balancing virtual appliance is supported by Citrix Hypervisor as documented in the Citrix Product Matrix (<https://www.citrix.com/support/product-lifecycle/product-matrix.html>).

Determine the best server 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, server. 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 server 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 recommenda-

tions, 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 Desktops environments.

To start a VM on the optimal server

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 server with the most stars.
 - **One of the servers with star ratings** listed under the Optimal Server command. Five stars indicate the most-recommended (optimal) server and five empty stars indicates the least-recommended server.

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 server

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 server with the most stars.
 - **One of the servers with star ratings** listed under the Optimal Server command. Five stars indicate the most-recommended (optimal) server and five empty stars indicates the least-recommended server.

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 master 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 master last when determining the order to fill hosts with VMs.

Optimization recommendations appear in the **WLB optimization** tab in XenCenter.

Optimization Recommendations

[View History...](#)

VM/Host	Operation	Reason
HA-prot-VM-7	Relocate from 'host17.domain4.bedford4.ctx4' to 'host16.domain4.be...	Consolidation
host17.domain4.bedford4.ctx4	Power off	Release Resource

Apply Recommendations

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. Citrix Hypervisor begins performing all the operations listed in the Operations column of the Optimization Recommendations section.

After you click Apply Recommendations, XenCenter automatically displays the Logs tab so you can see the progress of the VM migration.

Understand Workload Balancing recommendations under high availability

If you have Workload Balancing and Citrix Hypervisor 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 server 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 Print button is disabled when the report is generating.
	Print Layout enables you to display a preview of the report before you print.
	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 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. Citrix Hypervisor 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 Mini-

imum Dynamic Memory (MB), Current Maximum Dynamic Memory (MB), Current Assigned Memory (MB), and Average Assigned Memory (MB).

- **Dynamic memory allocation.** If you configured Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor Audit Log. The Audit Log is a Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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 [Citrix Hypervisor 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

Pool Audit Trail Granularity	Event Action	User Action
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
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

Pool Audit Trail Granularity	Event Action	User Action
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 history 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 server where the VM was originally hosted.
- **To Host:** The physical server 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

July 18, 2024

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 servers 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 servers.

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 server in the pool. For more information, see [Generate workload reports](#).

Notes:

- Workload Balancing is available for Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, see [Licensing](#). To upgrade, or to get a Citrix Hypervisor license, visit the [Citrix](#)

[website](#).

- Workload Balancing 8.2.2 and later are compatible with Citrix Hypervisor 8.2 Cumulative Update 1.
- The base OS of the Workload Balancing virtual appliance was originally derived from a general-purpose Linux distribution. The Workload Balancing virtual appliance is supported by Citrix Hypervisor as documented in the Citrix Product Matrix (<https://www.citrix.com/support/product-lifecycle/product-matrix.html>).

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 servers in a resource pool. The goal is to minimize CPU, memory, and network pressure for all servers. When Maximize Performance is your placement strategy, Workload Balancing recommends optimization when a server reaches the High threshold.

- **Maximize Density**

Workload Balancing attempts to minimize the number of physical servers 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 server. 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 servers in a resource pool. The goal is to minimize CPU, memory, and network pressure for all servers.
 - **Maximize Density**. Attempts to fit as many VMs as possible onto a physical server. The goal is to minimize the number of physical servers 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 servers in a resource pool. The goal is to minimize CPU, memory, and network pressure for all servers.
 - **Maximize Density.** Attempts to fit as many VMs as possible onto a physical server. The goal is to minimize the number of physical servers 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 servers on or off automatically. To power down servers 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 servers. However, to power down servers 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 servers 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 server 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 servers when doing so improves server performance.
 - Maximum Density mode: When **Automatically apply Power Management recommendations** is enabled, Workload Balancing automatically powers off servers when resource utilization drops below the Low threshold. That is, Workload Balancing powers off servers 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 servers that you want Workload Balancing to recommend powering on and off.

Note:

Selecting servers 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 servers 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 servers. In a Workload Balancing context, this term means powering servers in a pool on or off based on the total workload of the pool.

Configuring Workload Balancing power management on a server requires that:

- The hardware for the server has remote power on/off capabilities.
- The Host Power On feature is configured for the server. To configure the Host Power On feature for the server, see [Configure Host Power On feature](#).
- The server has been explicitly selected as a server to participate in Workload Balancing power management.

In addition, if you want Workload Balancing to power off servers automatically, configure Workload Balancing to do the following actions:

- Apply recommendations automatically
- Apply power management recommendations automatically

When a server 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 servers until it eliminates all excess capacity.
- If there isn't enough server capacity in the pool to shut down servers, Workload Balancing recommends leaving the servers on until the pool workload decreases enough.
- When you configure Workload Balancing to power off extra servers 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 servers automatically, Workload Balancing powers on servers when resource utilization on a server exceeds the High threshold.
- Workload Balancing never powers off servers 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 servers from the pool you want to participate in power management.

Configure Host Power On feature To configure the Host Power On feature for your server, follow these steps:

1. In Citrix Hypervisor Center, select your server and click **Properties**.
2. In the left pane, click **Power On**.
3. For the **Power On mode**, select **Dell Remote Access Controller (DRAC)**.
4. For the **Configuration options**, enter your server's DRAC IP address. This is the IP address of the BMC management port. For more information, see [DRAC Card How To Guide \[PDF\]](#).
5. After the Dell Remote Access Controller (DRAC) is configured, select your pool.
6. In the **Properties** pane of the pool, click the **WLB** tab.
7. In the **WLB** tab, click **Settings**.
8. In the left pane, click **Automation**.
9. 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 servers when doing so improves server performance.
 - Maximum Density mode: When **Automatically apply Power Management recommendations** is enabled, Workload Balancing automatically powers off servers when resource utilization drops below the Low threshold. That is, Workload Balancing powers off servers automatically during low usage periods.
10. For **Power Management**, select the name of the **Host Server** that you're currently configuring.

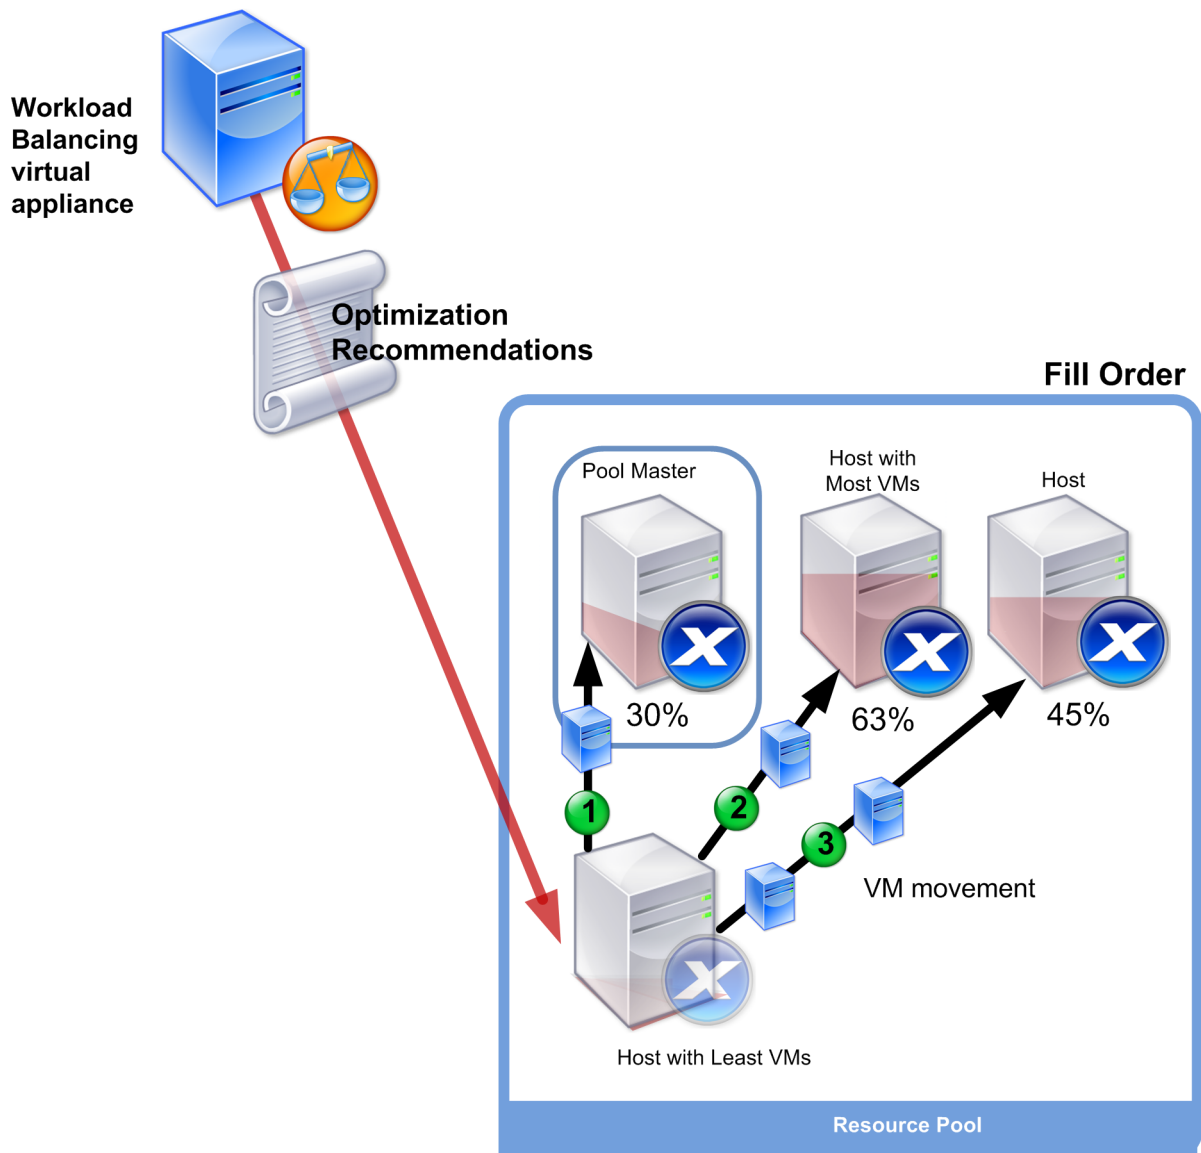
Understand power management behavior

Before Workload Balancing recommends powering servers on or off, it selects the servers to transfer VMs to. It does so in the following order:

1. Filling the pool master since it is the server that cannot be powered off.
2. Filling the server with the most VMs.
3. Filling subsequent servers according to which servers have the most VMs running.

When Workload Balancing fills the pool master, it does so assuming artificially low thresholds for the master. Workload Balancing uses these low thresholds as a buffer to prevent the pool master from being overloaded.

Workload Balancing fills servers 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 servers. If Workload Balancing cannot resolve the issue using this method, it attempts to power on a server. Workload Balancing determines which servers 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 servers 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 servers automatically or recommends doing so.

Important:

Workload Balancing only recommends powering on a server that Workload Balancing powered off.

Design environments for power management and VM consolidation

When you are planning a Citrix Hypervisor 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 pool.

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 servers from Workload Balancing.

Some servers might need to be always on. For more information, see [Exclude servers from recommendations](#).

Understand when Workload Balancing makes recommendations

Workload Balancing continuously evaluates the resource metrics of physical servers and VMs across the pools that it is managing against thresholds. Thresholds are preset values that function like boundaries that a server 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 servers it recommends function as the destination servers and in what order to make any optimizations. A destination server is the server where Workload Balancing recommends relocating one or more VMs.
4. Workload Balancing makes an optimization recommendation.

When evaluating servers 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 servers 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 servers' 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 server:

- 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 servers)
- 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 server 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 server 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 12:32 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 servers.
2. Determine placement recommendations: which servers would be suitable candidates for new VMs.

Note:

Workload Balancing only recommends migrating VMs that meet the Citrix Hypervisor criteria for live migration. One of these criteria is that the destination server must have the storage the VM requires. The destination server 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 server in the pool. It does so by monitoring on each server 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 server. If Workload Balancing can find another server with less CPU utilization, it recommends moving one or more VMs to another server.

2. If a resource's utilization on a server 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 servers 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 servers can support the VMs it wants to migrate off servers.

Workload Balancing makes this determination by calculating the projected effect on resource utilization of placing different combinations of VMs on servers. 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 server. The score indicates the suitability of a server as a home for more VMs.

To score server performance, Workload Balancing combines the following metrics:

- The current metrics of the server
- The metrics of the server from the last 30 minutes
- The metrics of the server from 24 hours ago

- The metrics of the VM.
5. After scoring servers and VMs, Workload Balancing attempts to build virtual models of what the servers look like with different combinations of VMs. Workload Balancing uses these models to determine the best server to place the VM.

In Maximum Performance mode, Workload Balancing uses metric weightings to determine what servers to optimize first and what VMs on those servers 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 servers and VMs to optimize according to the following criteria:

- a) What servers 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 servers as candidates for optimization and VMs as candidates for migration until predicted resource utilization on the server hosting the VM drops below the High threshold. Predicted resource utilization is the resource utilization that Workload Balancing forecasts a server has after Workload Balancing has added or removed a VM from the server.

Consolidation process in Maximum Density mode Workload Balancing determines whether to make a recommendation based on whether it can migrate a VM onto a server and still run that server 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 server, it evaluates whether the destination server is a suitable home for the VM.
3. Like in Maximum Performance mode, Workload Balancing scores the server to determine the suitability of a server as a home for new VMs.

Before Workload Balancing recommends consolidating VMs on fewer servers, it checks that resource utilization on those servers 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 servers.

4. After scoring servers and VMs, Workload Balancing attempts to build virtual models of what the servers look like with different combinations of VMs. It uses these models to determine the best server to place the VM.

5. Workload Balancing calculates the effect of adding VMs to a server until it forecasts that adding another VM causes a server resource to exceed the Critical threshold.
6. Workload Balancing recommendations always suggest filling the pool master first since it is the server that cannot be powered off. However, Workload Balancing applies a buffer to the pool master so that it cannot be over-allocated.
7. Workload Balancing continues to recommend migrating VMs on to servers until all remaining servers 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 servers 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 server 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 server with lower resource utilization.
 - **Maximum Density.** Workload Balancing doesn't recommend placing a VM on server when moving that VM to the server causes the server 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 servers are underutilized and makes an optimization recommendation to consolidate VMs on fewer servers. Workload Balancing continues to recommend moving VMs onto a server until the metric values for one of the server's resources reaches its High threshold.

However, after a VM is relocated, utilization of a resource on the VM's new server 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 server for the VMs. Workload Balancing continues to use this algorithm to recommend moving VMs until resource utilization on servers 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 server 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 server.

Workload Balancing uses these thresholds when making VM placement and pool-optimization recommendations. Workload Balancing strives to keep resource utilization on a server below the critical values set.

Tune metric weightings

How Workload Balancing uses metric weightings when determining which servers 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 servers according to which servers 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 servers 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 server 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 is in the same place for all resources.

Maximum Density mode

In Maximum Density mode, Workload Balancing only uses metric weightings when a server reaches the Critical threshold. At that point, Workload Balancing applies an algorithm similar to the algorithm for Maximum Performance until no servers 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 servers have resources exceeding their Critical thresholds, Workload Balancing verifies the importance you set for each resource. It uses this importance to determine which server to optimize first and which VMs on that server to relocate first.

For example, your pool contains server A and server B, which are in the following state:

- The CPU utilization on server A exceeds its Critical threshold and the metric weighting for CPU utilization is set to **More Important**.
- The memory utilization on server B exceeds its Critical threshold and the metric weighting for memory utilization is set to **Less Important**.

Workload Balancing recommends optimizing server 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 server A, Workload Balancing then begins recommending placements for VMs on that server. 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 server A, it makes optimization recommendations for server B. When it recommends placements for the VMs on server B, it does so by addressing CPU utilization first, since CPU utilization was assigned the highest weight. If there are more servers

that need optimization, Workload Balancing addresses the performance on those servers according to what server 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 be 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 servers from recommendations

When configuring Workload Balancing, you can specify that specific physical servers are excluded from Workload Balancing optimization and placement recommendations, including Start On placement recommendations.

Situations when you might want to exclude servers from recommendations include when:

- You want to run the pool in Maximum Density mode and consolidate and shut down servers, but you want to exclude specific servers from this behavior.
- You have two VM workloads that must always run on the same server. 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 server and you do not want VMs placed on the server.
- The performance of the workload is so critical that the cost of dedicated hardware is irrelevant.
- Specific servers are running high-priority workloads, and you do not want to use the HA feature to prioritize these VMs.
- The hardware in the server is not the optimum for the other workloads in the pool.

Regardless of whether you specify a fixed or scheduled optimization mode, excluded servers remain excluded even when the optimization mode changes. Therefore, if you only want to prevent Workload Balancing from shutting off a server automatically, consider disabling Power Management for that server instead. For more information, see [Optimize and manage power automatically](#).

When you exclude a server from recommendations, you are specifying for Workload Balancing not to manage that server at all. This configuration means that Workload Balancing doesn't make any optimization recommendations for an excluded server. In contrast, when you don't select a server to participate in Power Management, Workload Balancing manages the server, but doesn't make power management recommendations for it.

To exclude servers from Workload Balancing

Use this procedure to exclude a server in a pool that Workload Balancing is managing from power management, server 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 servers for which you do not want Workload Balancing to recommend alternate placements and optimizations.

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 servers) Workload Balancing requires before applying recommendations automatically

In general, only adjust the settings for these factors in the following cases:

- You have guidance from Citrix 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 servers 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 rec-

ommendation 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 server has stable enough performance over a prolonged period to accept a relocated VM without needing to move it off the server 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 server 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 server 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 Citrix 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 server or VM exhibits consistent behavior over time. Consistent behavior over time includes factors like whether a server continues to trigger recommendations and whether the same VMs on that server 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 server has stable resource utilization over a longer period. It reduces the potential for a recently moved VM to be moved off a server due to server 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 servers are being optimized. If you think that your servers 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 servers 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 server; however, they can be different VMs from the ones in the first recommendation.
- All destination servers 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 server. However, the recommendations do not have to follow each other immediately.
- The server 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 server 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 server A to server B
- Move VM3 from server A to server B
- Move VM5 from server A to server 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 server A to server B
- Move VM3 from server A to server C
- Move VM7 from server A to server 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 server A to server B
- Move VM3 from server A to server C
- Move VM5 from server A to server 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 server A to server B
- Move VM3 from server A to server B
- Move VM5 from server A to server C

Recommendation counts:

- Medium Aggressiveness Recommendation Count: 2
- Low Aggressiveness Recommendation Count: 1

Recommendation 5:

Proposed placements:

- Move VM1 from server A to server B
- Move VM3 from server A to server B
- Move VM5 from server A to server C

Recommendation counts:

- Medium Aggressiveness Recommendation Count: 3 (Apply)
- Low Aggressiveness Recommendation Count: 2

Recommendation 6:

Proposed placements:

- Move VM1 from server A to server B
- Move VM3 from server A to server B
- Move VM5 from server A to server 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 server.
 - 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 server 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

Via XenAPI, you can set the alert level for Workload Balancing alerts in XenCenter.

Complete the following steps:

1. Run the following command on the pool master 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 other 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 master 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 click the `Notifications` panel to view the alert.

Administer Workload Balancing

July 8, 2024

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.

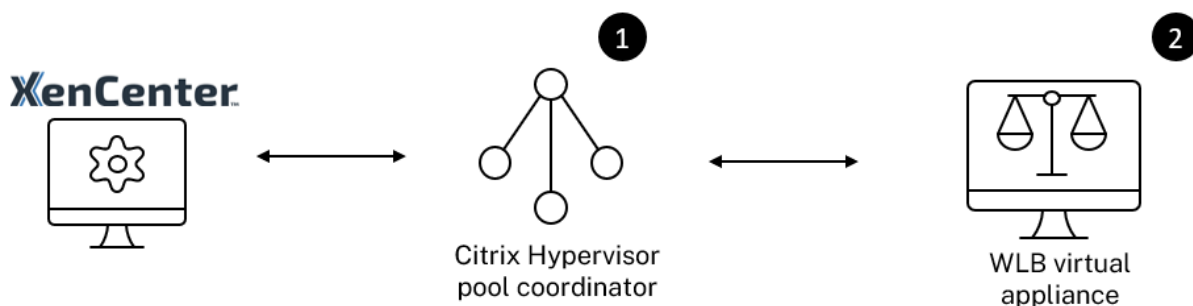
To connect a pool to your Workload Balancing virtual appliance, you need the following information:

- 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, Citrix Hypervisor 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. Citrix Hypervisor 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

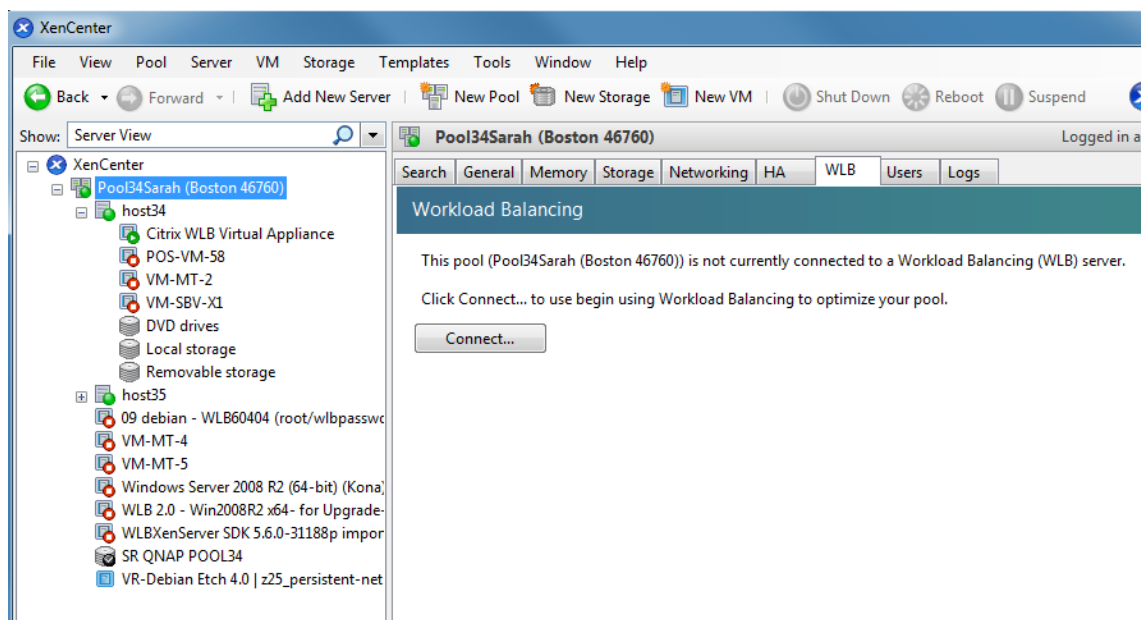
If you want to upload a different (trusted) certificate or configure certificate verification, note the following before connecting your pool to Workload Balancing:

- If you want Citrix Hypervisor 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.

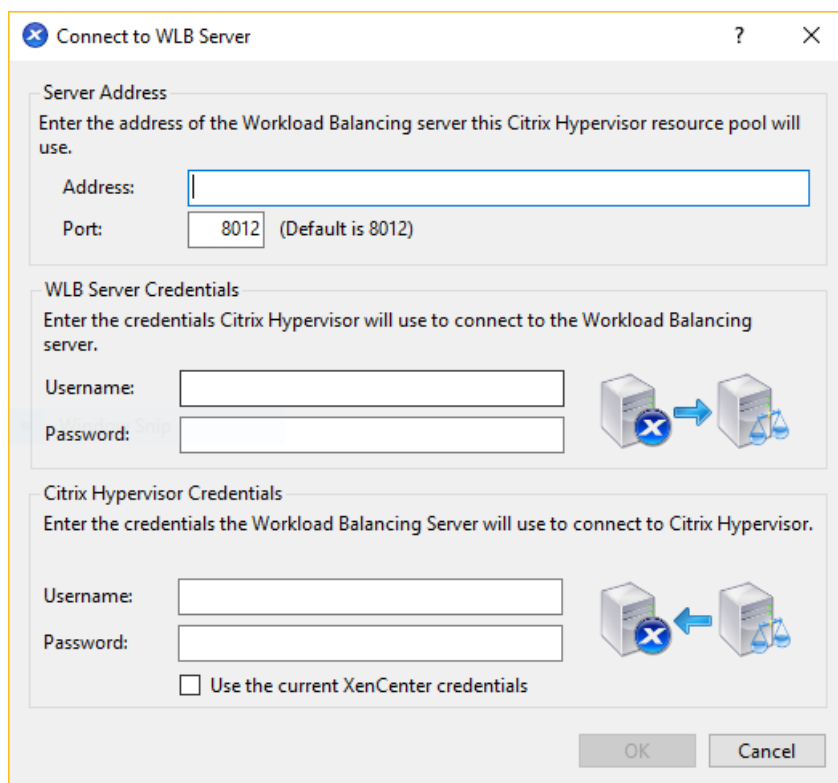
For more information, see [Certificates](#).

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.



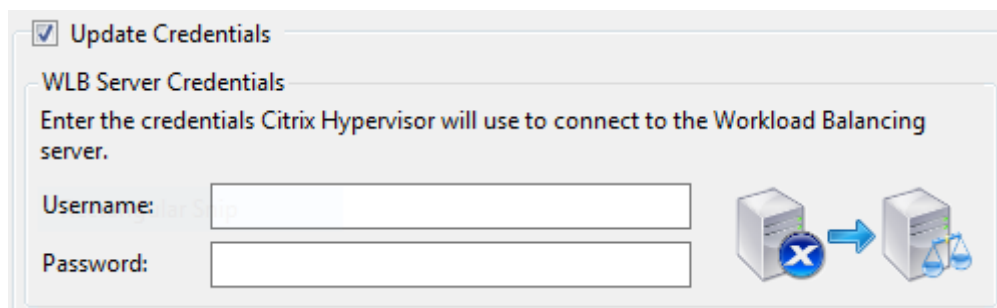
The dialog box titled "Connect to WLB Server" contains three sections: "Server Address", "WLB Server Credentials", and "Citrix Hypervisor Credentials". The "Server Address" section has an "Address" text box and a "Port" dropdown menu set to "8012" with a note "(Default is 8012)". The "WLB Server Credentials" section has "Username" and "Password" text boxes, a "Skip" button, and an icon showing a server connecting to a WLB server. The "Citrix Hypervisor Credentials" section has "Username" and "Password" text boxes, a "Skip" button, an icon showing a WLB server connecting to a Citrix Hypervisor server, and a checkbox labeled "Use the current XenCenter credentials". At the bottom are "OK" and "Cancel" buttons.

3. In the **Server Address** section, enter the following:

- In the **Address** box, type the IP address or FQDN of the Workload Balancing virtual appliance. For example, `WLB-appliance-computername.yourdomain.net`.
- (Optional) If you changed the Workload Balancing port during Workload Balancing Configuration, enter the port number in the **Port** box. Citrix Hypervisor uses this port to communicate with Workload Balancing.

By default, Citrix Hypervisor connects to Workload Balancing on port 8012.

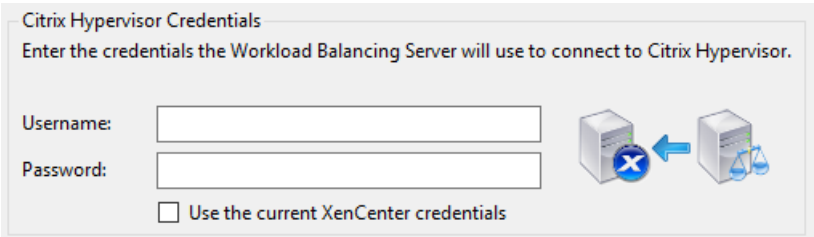
4. In the **WLB Server Credentials** section, enter the user name and password that the pool uses to connect to the Workload Balancing virtual appliance.



The "Update Credentials" dialog box has a checked checkbox at the top. It contains a "WLB Server Credentials" section with "Username" and "Password" text boxes, a "Skip" button, and an icon showing a server connecting to a WLB server.

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 **Citrix Hypervisor Credentials** section, enter the user name and password for the pool you are configuring. Workload Balancing uses these credentials to connect to the servers in that 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 showing 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 Citrix Hypervisor, 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](#).

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

Permission	Minimum Required Role
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 servers 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 Citrix Hypervisor 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 Citrix Hypervisor uses to connect to Workload Balancing.

By default, Citrix Hypervisor 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 Citrix Hypervisor uses to connect to the Workload Balancing server.

8. In the **Citrix Hypervisor Credentials** section, enter the user name and password for the pool you are configuring (typically the password for the pool master). Workload Balancing uses these credentials to connect to the computers running Citrix Hypervisor in that pool.
9. In the **Citrix Hypervisor Credentials** section, enter the user name and password for the pool you are configuring. Workload Balancing uses these credentials to connect to the computers running Citrix Hypervisor in that pool.

To use the credentials with which you are currently logged into Citrix Hypervisor, select **Use the current XenCenter credentials**.

Change the Workload Balancing IP address

To change the Workload Balancing IP address, complete the following:

1. To view the current Workload Balancing IP address, run the `ifconfig` command on the virtual appliance.
2. Open the `/etc/sysconfig/network-scripts/ifcfg-eth0` file by using an editing tool like vi.
3. To change the protocol from dhcp to static, change `BOOTPROTO=dhcp` to `BOOTPROTO=static`.
4. At the bottom of the file, set the IP address, netmask, gateway, and DNS addresses. For example:

```
1 IPADDR=192.168.1.100
2 NETMASK=255.255.255.0
3 GATEWAY=192.168.1.1
4 DNS1=1.1.1.1
5 DNS2=8.8.8.8
```

Note:

Add as many DNS entries as you need.

5. Save and close the file.
6. For the changes to take effect, you must restart the networking system by running `systemctl restart network`.
7. Once the networking system has restarted, run the `ifconfig` command again to view the new Workload Balancing IP address.
8. 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 affect 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 8.2.2 and earlier This procedure explains how to extend the virtual disk of the Workload Balancing virtual appliance for Workload Balancing versions 8.2.2 and earlier.

Warning:

We recommend taking a snapshot of your data before performing this procedure. Incorrectly performing these steps can result in corrupting the Workload Balancing virtual appliance.

1. Shut down the Workload Balancing virtual appliance.
2. In the XenCenter resource pane, select the Workload Balancing virtual appliance.
3. Click the **Storage** tab.
4. Select the `vdi_xvda` disk, and click the **Properties** button.
5. In the `vdi_xvda` **Properties** dialog, select **Size and Location**.
6. Increase the disk size as needed, and click **OK**.
7. Start the Workload Balancing virtual appliance and log in to it.
8. Run the following command on the Workload Balancing virtual appliance:

```
1 resize2fs /dev/xvda
```

9. Run the `df -h` command to confirm the new disk size.

Installing `resize2fs` If the `resize2fs` tool is not installed on the Workload Balancing virtual appliance, you can install it by using the following steps.

If you are connected to the internet, run the following command on the Workload Balancing virtual appliance:

```
1 yum install -y --enablerepo=base,updates --disablerepo=citrix-*  
  e2fsprogs
```

If there is no internet access:

1. Download the following from https://centos.pkgs.org/7/centos-x86_64/.
 - `libss-1.42.9-7.el7.i686.rpm`
 - `e2fsprogs-libs-1.42.9-7.el7.x86_64.rpm`
 - `e2fsprogs-1.42.9-7.el7.x86_64.rpm`
2. Upload them to Workload Balancing VM using SCP or any other suitable tool.
3. Run the following command from Workload Balancing virtual appliance:

```
1 rpm -ivh libss-*.rpm e2fsprogs-*.rpm
```

The tool `resize2fs` is now installed.

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. Citrix cannot guarantee that problems resulting from the incorrect use of the Registry Editor can be solved. Be sure to back up the registry before you edit it and shut down the virtual appliance before performing these steps. 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 3149KB 2099KB grub         bios_grub
 2      3149KB 19.9MB 16.8MB ext4         grubConfig
 3      19.9MB 32.2GB 32.2GB               rootfs

(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
```

- 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.

- 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'
```

- 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/xvda3		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 +3081 /dev/centos/root
  Size of logical volume centos/root changed from <6.67 GiB (1707 extents) to 18
.70 GiB (4788 extents).
  Logical volume centos/root successfully resized.
-bash-4.2# xfs_growfs /dev/centos/root
meta-data=/dev/mapper/centos-root isize=256    agcount=4, agsize=436992 blks
         =                       sectsz=512    attr=2, projid32bit=1
         =                       crc=0        finobt=0 spinodes=0
data     =                       bsize=4096    blocks=1747968, imaxpct=25
         =                       sunit=0      swidth=0 blks
naming   =version 2              bsize=4096    ascii-ci=0 ftype=0
log      =internal              bsize=4096    blocks=2560, version=2
         =                       sectsz=512    sunit=0 blks, lazy-count=1
realtime =none                  extsz=4096    blocks=0, rtextents=0
data blocks changed from 1747968 to 4902912
```

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 Citrix 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 server in maintenance mode, Citrix Hypervisor migrates the VMs running on that server to their optimal servers when available. Citrix Hypervisor uses Workload Balancing recommendations that are based on performance data, your placement strategy, and performance thresholds to select the optimal server.

If an optimal server 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 server 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 server in the same pool. Then, if you reenter the **Enter Maintenance Mode** dialog box, Workload Balancing might be able to list a server that is a suitable candidate for migration.

Note:

When you take a server off-line 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 server are automatically migrated to the optimal server 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 server and select **Exit Maintenance Mode**.

When you remove a server from maintenance mode, Citrix Hypervisor automatically restores that server’s original VMs to that server.

Remove the Workload Balancing virtual appliance

To remove the Workload Balancing virtual appliance, we recommend 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 firewall configured in it. Before you can access the database, you must add the postgresQL server port to the iptables.

1. From the Workload Balancing virtual appliance console, run the following command:

```
1 iptables -A INPUT -i eth0 -p tcp -m tcp --dport 5432 -m \
2 state --state NEW,ESTABLISHED -j ACCEPT
```

2. (Optional) To make this configuration persist after the virtual appliance is rebooted, run the following command:

```
1 iptables-save > /etc/sysconfig/potables
```

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 1,024 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 2,048 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/9.0/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 service postgresql-9.0 restart
```

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/9.0/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 service postgresql-9.0 restart
```

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:

To avoid a loss of Workload Balancing:

- During the maintenance window, the Workload Balancing server restarts. Ensure that you do not restart your VMs at the same time.
- At other times, when restarting all VMs in your pool, do not restart the Workload Balancing server.

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 <https://www.citrix.com/downloads/citrix-hypervisor/product-software/>.

Workload Balancing commands

This section provides a reference for the Workload Balancing commands. You can perform these commands from the Citrix Hypervisor server or console to control Workload Balancing or configure Work-

load Balancing settings on the Citrix Hypervisor server. 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 Citrix. 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, Citrix recommends 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, Citrix recommends using caution.

For all other settings in the `wlb.conf` file, Citrix currently recommends leaving them at their default, unless Citrix instructed you 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 Citrix Technical Support or when troubleshooting.

Logging Option	Trace Flag	Benefit or Purpose
Analysis Engine Trace	<code>AnalEngTrace</code>	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	<code>DatabaseTrace</code>	Logs details about database reads/writes. However, leaving this trace on increases the log file size quickly.
Data Collection Trace	<code>DataCollectionTrace</code>	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	<code>DataCompactionTrace</code>	Logs details about how many milliseconds it took to compact the metric data.
Data Event Trace	<code>DataEventTrace</code>	This trace provides details about events Workload Balancing catches from Citrix Hypervisor.
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.)

Logging Option	Trace Flag	Benefit or Purpose
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 servers 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.
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

January 9, 2023

Citrix Hypervisor and Workload Balancing communicate over HTTPS. During Workload Balancing Con-

figuration, the wizard automatically creates a self-signed test certificate. This self-signed test certificate lets Workload Balancing establish a TLS connection to Citrix Hypervisor. By default, Workload Balancing creates this TLS connection with Citrix Hypervisor automatically. You do not need to perform any certificate configurations during or after configuration for Workload Balancing to create this TLS connection.

Note:

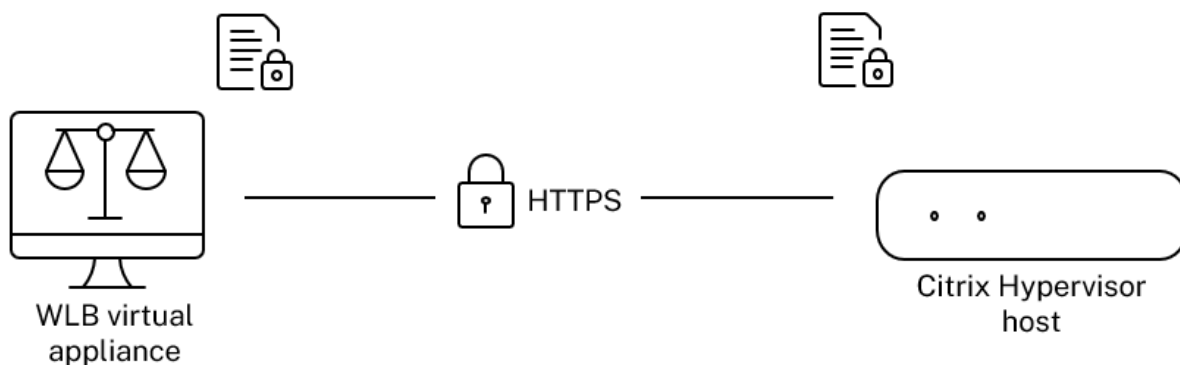
The self-signed certificate is a placeholder to facilitate HTTPS communication and is not from a trusted certificate authority. For added security, we recommend using a certificate signed from a trusted certificate authority.

To use a certificate from another certificate authority, such as a signed one from a commercial authority, you must configure Workload Balancing and Citrix Hypervisor to use it.

By default, Citrix Hypervisor does not validate the identity of the certificate before it establishes connection to Workload Balancing. To configure Citrix Hypervisor to check for a specific certificate, export the root certificate that was used to sign the certificate. Copy the certificate to Citrix Hypervisor and configure Citrix Hypervisor to check for it when a connection to Workload Balancing is made. Citrix Hypervisor acts as the client in this scenario and Workload Balancing acts as the server.

Depending on your security goals, you can either:

- [Configure Citrix Hypervisor to verify the self-signed certificate.](#)
- [Configure Citrix Hypervisor 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, Citrix Hypervisor doesn't allow the connection.

For certificate verification to succeed, you must store the certificates in the specific locations in which Citrix Hypervisor expects to find the certificates.

Configure Citrix Hypervisor to verify the self-signed certificate

You can configure Citrix Hypervisor to verify that the Citrix Workload Balancing self-signed certificate is authentic before Citrix Hypervisor permits Workload Balancing to connect.

Important:

To verify the Citrix 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 Citrix Hypervisor 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 master. The Citrix Workload Balancing self-signed certificate is stored at `/etc/ssl/certs/server.pem`. Run the following command on the pool master:

```
1 scp root@<wlb-ip>:/etc/ssl/certs/server.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=server.pem
```

5. Verify the certificate was installed correctly by running the following command on the pool master:

```
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 master to all servers in the pool, running the following command on the pool master:

```
1 xe pool-certificate-sync
```

Running the `pool-certificate-sync` command on the master synchronizes the certificate and certificate revocation lists on all the pool servers with the master. This action ensures all servers 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 Citrix Hypervisor to verify the certificate before connecting to the Workload Balancing virtual appliance. Run the following command on the pool master:

```
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 servers in the pool, run the `pool-certificate-list` command on those servers.
 - b) To test if Citrix Hypervisor 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 Citrix Hypervisor to verify a certificate-authority certificate

You can configure Citrix Hypervisor to verify a certificate signed by a trusted certificate authority.

For trusted authority certificates, Citrix Hypervisor 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 Citrix Hypervisor pool master.
- Citrix Hypervisor can resolve the Workload Balancing host name. (For example, you can try pinging the Workload Balancing FQDN from the Citrix Hypervisor console for the pool master.)

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 Workload Balancing use 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 killall stunnel
2 stunnel
```

Import the certificate chain into the pool

After you obtain the certificates, import them onto the Citrix Hypervisor pool master. Synchronize the servers in the pool to use those certificates. Then you can configure Citrix Hypervisor to check the certificate identity and validity each time Workload Balancing connects to a server.

1. Copy the signed certificate, root certificate and, if the certificate authority has one, the intermediate certificate from the certificate authority onto the Citrix Hypervisor pool master.
2. Install the root certificate on the pool master:

```
1 xe pool-certificate-install filename=root_ca.pem
```

3. If applicable, install the intermediate certificate on the pool master:

```
1 xe pool-certificate-install filename=intermediate_ca.pem
```

4. Verify both the certificates installed correctly by running this command on the pool master:

```
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 master to all servers in the pool:

```
1 xe pool-certificate-sync
```

Running the `pool-certificate-sync` command on the master synchronizes the certificates and certificate revocation lists on all the pool servers with the pool master. This action ensures all servers in the pool use the same certificates.

6. Instruct Citrix Hypervisor to verify a certificate before connecting to the Workload Balancing virtual appliance. Run the following command on the pool master:

```
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** 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** dialog exactly as it appears in the certificate's Common Name. Enter the FQDN to ensure that the Common Name matches the name that Citrix Hypervisor uses to connect.

Troubleshooting

- If the pool cannot connect to Workload Balancing after configuring certificate verification, check to see if the pool can connect if you turn certificate verification off. You can use the command `xe pool-param-set wlb-verify-cert=false uuid=uuid_of_pool` to disable certificate verification. If it can connect with verification off, the issue is in your certificate configuration. If it cannot connect, the issue is in either your Workload Balancing credentials or your network connection.
- 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](#).

Troubleshoot Workload Balancing

June 11, 2024

While Workload Balancing usually runs smoothly, this series of sections provides guidance in case you encounter issues.

Notes:

- Workload Balancing is available for Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, see [Licensing](#). To upgrade, or to get a Citrix Hypervisor license, visit the [Citrix website](#).
- Workload Balancing 8.2.2 and later are compatible with Citrix Hypervisor 8.2 Cumulative Update 1.

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 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 **Logs** tab 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 analyse 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 **Logs** tab 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 Workload Balancing virtual appliance imported and was configured correctly and 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:** Citrix Hypervisor 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`.
 - **Citrix Hypervisor Credentials:** This account is used by the Workload Balancing virtual appliance to connect to the Citrix Hypervisor pool. This account is created on the Citrix Hypervisor pool master and has the `pool-admin` or `pool-operator` role.
- You can enter a host name in the **Address** box, but it must be the fully qualified domain name (FQDN) of the Workload Balancing virtual appliance. Do not enter the host name of the physical server hosting the appliance. If you are having trouble entering a computer name, try using the Workload Balancing appliance's IP address instead.
- Verify that the host is using the correct DNS server and the Citrix Hypervisor server can contact Workload Balancing server using its FQDN. To do this check, ping the Workload Balancing appliance using its FQDN from the Citrix Hypervisor server. For example, enter the following in the Citrix Hypervisor server console:

```
1 ping wlb-vpx-1.mydomain.net
```

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.

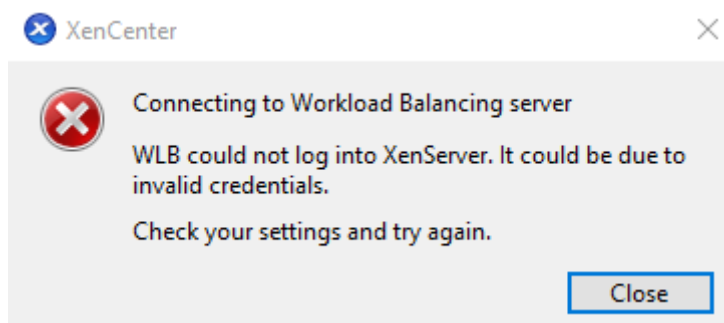
Likewise, the port Citrix Hypervisor uses to contact Workload Balancing (8012 by default), must match the port number specified when you ran the Workload Balancing Configuration wizard.

Workload Balancing connection errors

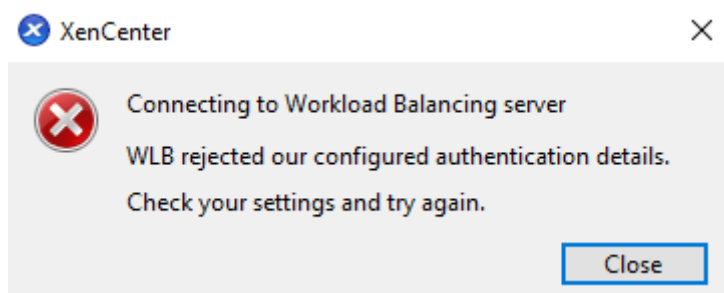
If you receive a connection error after configuring and connecting to Workload Balancing, the credentials might no longer be valid. To isolate this issue:

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 server 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 **Citrix Hypervisor 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

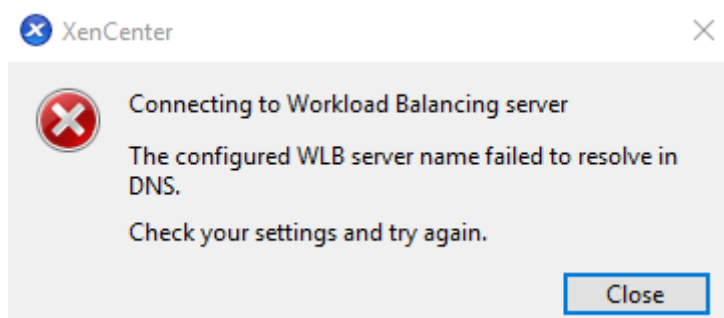
This means that there is a problem with the credentials entered in the **WLB Server Credentials** field in the **Connect to WLB Server** dialog box when attempting to connect to the Workload Balancing virtual appliance (either the username or the password are incorrect). However, it can also mean that 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`

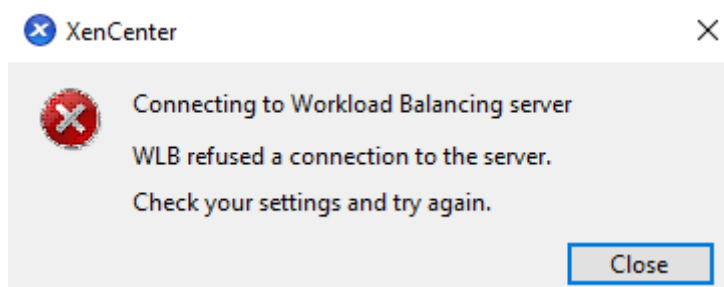
Scenario 3



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 <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):

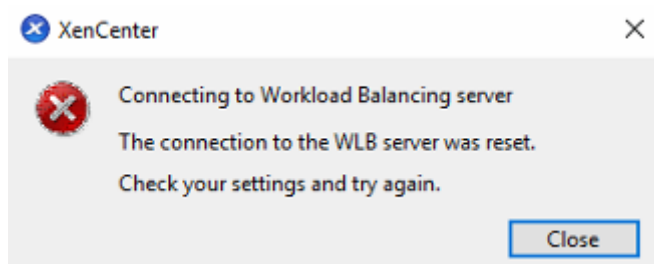
1. Run the following command in the WLB appliance CLI:

```
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  dwmdatacolsvc.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 Citrix Hypervisor 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. Temporarily disable HA, if possible, 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.

XenServer Conversion Manager

June 11, 2024

Use the XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager) virtual appliance to migrate your VMs running on VMware ESXi/vCenter to Citrix Hypervisor quickly and efficiently. You can batch convert multiple VMs in parallel at the same time.

As part of the migration, XenCenter helps you prepare the VMs for networking and storage connectivity. After converting your VMs to a Citrix Hypervisor environment, they're almost ready to run. Complete the post-conversion tasks which include installing XenServer VM Tools (formerly Citrix VM Tools) for Windows on your Windows VMs. For Linux VMs, the XenServer Conversion Manager automatically installs XenServer VM Tools for Linux during the conversion process.

After converting your VMs, the Conversion Manager automatically shuts down by itself, saving resources on the host.

Note:

In Citrix Hypervisor 8.0 and earlier, a separate Conversion Manager console is provided. From Citrix Hypervisor 8.1, this capability is integrated into XenCenter.

Overview

Citrix Hypervisor allows you to:

- Batch convert multiple VMware ESXi/vCenter VMs in parallel using one simple wizard
- Map network settings between VMware and Citrix Hypervisor so your converted VMs can be up and running with the proper network settings
- Select a storage location where you would like your new Citrix Hypervisor VMs to run

Notes:

- XenCenter does not remove or change your existing VMware environment. VMs are duplicated onto your Citrix Hypervisor environment and not removed from VMware.

- XenServer Conversion Manager virtual appliance supports converting VMware ESXi/vCenter VMs with different storage such as thin provisioning, thick provisioning, IDE, and SCSI.
- XenServer 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.
- XenServer Conversion Manager virtual appliance cannot convert VMware ESXi/vCenter VMs with four or more disks into Citrix Hypervisor VMs. Your VMware ESXi/vCenter VMs must have three or fewer disks.
- XenServer Conversion Manager virtual appliance is available for Citrix Hypervisor Premium Edition customers or customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, see [Licensing](#). To upgrade, or to get a Citrix Hypervisor 8.2 license, visit the [Citrix website](#).

Understand Citrix Hypervisor

Before you can convert your environment, it is suggested that you become familiar with Citrix Hypervisor concepts. For more information, see [Technical overview](#).

To successfully convert VMware virtual machines to Citrix Hypervisor, perform the following tasks:

- Set up a basic Citrix Hypervisor environment, including installing Citrix Hypervisor. For more information, see [Quick start](#) and [Install](#).
- Create a network in Citrix Hypervisor, 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 Citrix Hypervisor terminology The following table lists the approximate Citrix Hypervisor equivalent for common VMware features, concepts, and components:

VMware Term	Citrix Hypervisor Equivalent
VMware vSphere Client	XenCenter (the management console for Citrix Hypervisor)
Cluster / Resource Pool	Resource Pool
Data Store	Storage Repository
vMotion	Live migration

VMware Term	Citrix Hypervisor Equivalent
Distributed Resource Scheduling (DRS)	Workload Balancing
High Availability (HA)	High Availability (HA)
vCenter Converter	XenServer Conversion Manager
Role Based Access Control (RBAC)	Role Based Access Control (RBAC)

Conversion overview

XenCenter and XenServer Conversion Manager virtual appliance create a copy of each targeted VM. After converting the targeted VM to a Citrix Hypervisor VM with comparable networking and storage connectivity, XenCenter imports the VM into your Citrix Hypervisor 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. XenServer Conversion Manager virtual appliance does not support converting a running VM using memory copied from vSphere to Citrix Hypervisor.

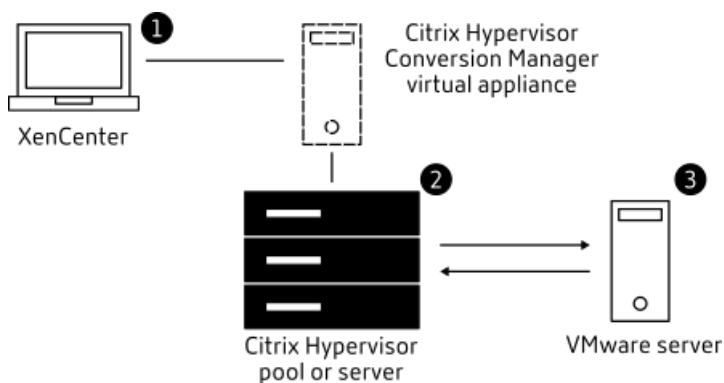
Also, before converting, ensure that a network and a storage controller exist in your VMware VM.

The conversion process requires four items:

- **XenCenter** - the Citrix Hypervisor 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 Citrix Hypervisor and the XenServer Conversion Manager virtual appliance.
- **XenServer Conversion Manager virtual appliance** - a pre-packaged VM you import into the Citrix Hypervisor host or pool where you want to run the converted VMs. The virtual appliance converts the copies of the VMware ESXi/vCenter VMs into Citrix Hypervisor virtual machine format. After conversion, it imports these copies into the Citrix Hypervisor pool or host.
- **Citrix Hypervisor standalone host or pool** - the Citrix Hypervisor environment where you want to run the converted VMs.
- **VMware server** - XenServer 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 XenServer Conversion Manager Virtual Appliance makes a copy of these VMs and converts them to Citrix Hypervisor virtual-machine format.

The following illustration shows the relationships between these components:



This illustration shows:

1. How XenCenter communicates with XenServer Conversion Manager virtual appliance.
2. How the XenServer Conversion Manager virtual appliance authenticates with the VMware server.
3. How the VMware server responds to the XenServer Conversion Manager virtual appliance during conversion.

The VMware server communicates with the XenServer 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 XenServer Conversion Manager virtual appliance and start to convert VMs in just a few easy steps:

1. Download the XenServer Conversion Manager virtual appliance from the [Citrix Hypervisor 8.2 Premium Edition page](#).
2. Import the XenServer Conversion Manager virtual appliance into Citrix Hypervisor using XenCenter.
3. Configure the XenServer Conversion Manager virtual appliance by using XenCenter.
4. From XenCenter, launch the conversion wizard and start to convert VMs.
5. Complete the post-conversion tasks which include installing XenServer VM Tools (formerly Citrix VM Tools) for Windows on your Windows VMs. For Linux VMs, the XenServer Conversion Manager automatically installs XenServer VM Tools for Linux during the conversion process.

After converting your VMs, the Conversion Manager automatically shuts down by itself, saving resources on the host. For more information on how to convert VMware ESXi/vCenter VMs, see [Get started with XenServer Conversion Manager](#).

What's new in XenServer Conversion Manager

May 10, 2024

The latest version of the XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager) virtual appliance is version 8.3.1. You can download this version of the XenServer Conversion Manager virtual appliance from the [Citrix Hypervisor Downloads page](#).

If you already have a previous version of the XenServer Conversion Manager virtual appliance installed and wish 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.

What's new in 8.3.1

Released Feb 01, 2024

This update includes the following improvement:

- You can now batch convert multiple VMware ESXi/vCenter VMs in parallel at the same time.

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.

Known issues in 8.3.1

- Windows VMs converted from VMware to XenServer using the XenServer Conversion Manager are incorrectly detected as Linux VMs.

To work around this issue, download and install XenServer VM Tools for Windows (available to download from the [XenServer downloads page](#)) and then reboot your VM.

Get started with XenServer Conversion Manager

July 8, 2024

Using the XenServer Conversion Manager (formerly Citrix Hypervisor Conversion Manager), you can easily convert your VMware ESXi/vCenter virtual machines (VMs) to Citrix Hypervisor in just a few steps:

1. [Prepare your Citrix Hypervisor environment and review the prerequisite information.](#)
2. [Import and configure the XenServer Conversion Manager virtual appliance by using XenCenter.](#)

Note:

If you already have a previous version of the XenServer Conversion Manager virtual appliance installed and wish 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 Citrix Hypervisor.](#)
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 Citrix Hypervisor standalone 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 Citrix Hypervisor?
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 Citrix Hypervisor 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 Citrix Hypervisor host (if the cluster is not on the same network as the Citrix Hypervisor host).

- Configure a network to connect to the storage array. That is, if you use IP-based storage, create a Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor and create a pool

Before you can convert VMware ESXi/vCenter VMs, ensure that you create a Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor, you can learn about Citrix Hypervisor basics, including basic installation and configuration, by reading [Quick start](#).

Citrix Hypervisor environment considerations

Before installing Citrix Hypervisor and importing the virtual appliance, consider the following factors that might change your conversion strategy:

Selecting the host where you want to run the XenServer Conversion Manager virtual appliance.

Import the virtual appliance into the stand-alone host or into 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 XenServer Conversion Manager in a pool at a time.

The storage configured for the pool or host where you want to run the converted VMs must meet specific requirements. If you want to run your newly converted VMs in a pool, their virtual disks must be stored on shared storage. However, if the converted VMs run on a single standalone host (not a pool), their virtual disks can use local storage.

If you want to run the converted VMs in a pool, ensure that you add the shared storage to the pool by creating a storage repository.

Guest operating systems supported for conversion:

You can convert VMware ESXi/vCenter VMs running the following Windows guest operating systems:

- Windows 10 (32-bit) Enterprise edition [Latest tested version is 22H2]
 - Only boot from BIOS mode is supported
- Windows 10 (64-bit) Enterprise edition [Latest tested version is 22H2]
- Windows Server 2016 (64-bit) Standard (Desktop) edition
- Windows Server 2019 (64-bit) Standard (Desktop) edition
- Windows Server 2022 (64-bit) Standard (Desktop) edition

Note:

Only the listed Windows SKUs are supported for conversion.

The following Linux operating systems are also supported:

- Red Hat Enterprise Linux 7.9 (64-bit) with the following configuration:
 - File system: EXT3 or EXT4
 - Boot partition type: btrfs, lvm, or plain
- Red Hat Enterprise Linux 8.x (64-bit) with the following configuration:
 - File system: EXT3 or EXT4
 - Boot partition type: lvm or plain
- Ubuntu 20.04 with the following configuration:
 - File system: EXT3 or EXT4
 - Boot partition type: lvm or regular

For more information about the guest operating systems supported by Citrix Hypervisor, see [Guest operating system support](#).

Meet networking requirements To convert VMware ESXi/vCenter VMs, the XenServer 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 Citrix Hypervisor pool, add the network to Citrix Hypervisor before conversion.

Note:

- The time it takes for your VMs to be converted depends on the physical distance between your VMware and Citrix Hypervisor 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 XenServer 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. For more information, see VMware’s article [Improving transfer speed of task with library items](#).

Map your existing network configuration XenServer 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 Citrix Hypervisor. For example, XenServer Conversion Manager virtual appliance will:

- Preserve virtual MAC addresses on the VMware ESXi/vCenter VMs and reuse them in the resulting Citrix Hypervisor 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. XenServer Conversion Manager virtual appliance can map VMware networks onto Citrix Hypervisor 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 Citrix Hypervisor ‘Network 0’, any VMware VM that had a virtual adapter connected to ‘Virtual Network 4’ is connected to ‘Network 0’ after conversion. XenServer 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 Citrix Hyper-

visor networks. However, if you prefer, you can change the networks the VMs use, reduce, or consolidate the number of networks in your new Citrix Hypervisor configuration.

To gain the maximum benefit from these features, Citrix recommends the following:

- Before installing Citrix Hypervisor, 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 Citrix Hypervisor pool can see the networks that you would like to be detected. Specifically, plug the Citrix Hypervisor hosts into switch ports that can access the same networks as the VMware cluster.

Though it is easier to plug the Citrix Hypervisor 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 Citrix Hypervisor NIC into a different physical network.

Prepare for the XenServer 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. XenServer Conversion Manager virtual appliance uses this connection for conversion traffic between the Citrix Hypervisor host and the VMware server.

To create this network connection, you must perform two tasks:

- When you import the XenServer 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 Citrix Hypervisor to the Citrix Hypervisor host where you want to run the converted VMs.

By default, when you import the XenServer Conversion Manager virtual appliance, XenCenter creates one virtual network interface associated with Network 0 and NIC0 (eth0). 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 XenServer conversion manager, you might see multiple network interfaces (eth0 and eth1). eth0 attaches to the host's internal network which is used to communicate with the local dom0. eth1 attaches to the routable network which is used to communicate with XenCenter.

To add a network to Citrix Hypervisor:

1. In the **Resource** pane in XenCenter, select the pool where you would like to run XenServer 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 Citrix Hypervisor 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**.

Meet storage requirements Before you convert batches of VMware ESXi/vCenter VMs, consider your storage requirements. Converted VM disks are stored on a Citrix Hypervisor 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. For converted machines that only run on a standalone host, you can specify either local or shared storage as the location for the converted virtual disks. For converted machines running in pools, you can only specify shared storage.

To create a storage repository:

1. In the **Resource** pane in XenCenter, select the pool where you intend to run the XenServer 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.

Citrix Hypervisor requirements You can run VMs converted with this release of XenServer Conversion Manager on the following versions of Citrix Hypervisor:

- Citrix Hypervisor 8.2 Cumulative Update 1

VMware requirements XenServer Conversion Manager virtual appliance can convert VMware ESXi/vCenter VMs from the following versions of VMware:

- vCenter Server 6.7.x, 7.x, and 8.x
- vSphere 6.7.x, 7.x, and 8.x
- ESXi 6.7.x, 7.x, and 8.x

Note:

XenServer Conversion Manager virtual appliance cannot convert VMware ESXi/vCenter VMs with four or more disks into Citrix Hypervisor VMs. Your VMware ESXi/vCenter VMs must have three or fewer disks.

Your VMware ESXi/vCenter VMs must also have a network and a storage controller configured.

Prepare to import the virtual appliance Before importing the virtual appliance, note the following information and make the appropriate changes to your environment, as applicable.

Download the virtual appliance The XenServer Conversion Manager virtual appliance is packaged in XVA format. You can download the virtual appliance from the [Citrix Hypervisor 8.2 Premium Edition 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:

XenServer Conversion Manager virtual appliance is available for Citrix Hypervisor Premium Edition customers or those customers who have access to Citrix Hypervisor through their Citrix Virtual Apps and Desktops entitlement or Citrix DaaS entitlement. For more information about Citrix Hypervisor licensing, See [Licensing](#). To upgrade, or to get a Citrix Hypervisor 8.2 license, visit the Citrix website.

Virtual appliance prerequisites The XenServer Conversion Manager virtual appliance requires a minimum of:

- Citrix Hypervisor 8.2 Cumulative Update 1
- Disk space: 30 GB of disk space
- Memory: 6 GB
- Virtual CPU allocation: 2 vCPU

Import and configure the virtual appliance

The XenServer Conversion Manager virtual appliance is a single pre-installed VM designed to run on a Citrix Hypervisor host. Before importing it, review the prerequisite information and considerations in the section called *Preparing to import the virtual appliance*.

Import the virtual appliance into Citrix Hypervisor

To import the XenServer 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 XenServer 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 Citrix Hypervisor attempts to start the VM on that host, before trying other hosts. If you select a host, the XenServer 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 XenServer 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 Citrix Hypervisor 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 Citrix Hypervisor**.

Warning:

Do NOT configure NIC0 to your customer network. Assign NIC0 only to "Host internal management network."

6. Leave the **Start VM after import** check box enabled, and click **Finish** to import the virtual appliance.
7. After importing the `.xva` file, the XenServer Conversion Manager virtual appliance appears in the **Resources** pane in XenCenter.

Configure the XenServer Conversion Manager virtual appliance

Before you can use the XenServer Conversion Manager virtual appliance to convert VMware ESXi/vCenter VMs, configure it using the XenCenter **Console** tab:

1. After importing the XenServer 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 XenServer Conversion Manager virtual appliance. Citrix recommends selecting a strong password.
4. Enter a host name for the XenServer 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 XenServer 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 XenServer 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 XenServer Conversion Manager virtual appliance.

If you do not configure the XenServer Conversion Manager correctly, you might have to reconfigure the virtual appliance. For more information, see [Reconfigure the XenServer Conversion Manager virtual appliance](#).

Reconfigure the XenServer Conversion Manager virtual appliance If you need to change the configuration of the XenServer Conversion Manager after installing the virtual appliance, you can follow these steps to reconfigure it:

1. Log in to the XenServer Conversion Manager virtual appliance by using SSH or the XenCenter **Console** tab as 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 XenCenter **Console** tab.

Convert VMware ESXi/vCenter VMs

When you convert VMware ESXi/vCenter VMs, they are imported into the Citrix Hypervisor pool or standalone host where you are running the XenServer 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 is true:

- You have the credentials for the Citrix Hypervisor pool (or standalone host). Either the root account credentials or a Role-Based Access Control (RBAC) account with the Pool Admin role configured is acceptable.
- You have the credentials for the VMware server containing the VMs you want to convert. The conversion procedure requires you connect the XenServer Conversion Manager Console to the VMware server.
- The VMware virtual machines to convert are powered off.
- The VMware virtual machines to convert have a network and a storage controller configured.
- The Citrix Hypervisor pool (or host) that run 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.
- The virtual disks of the VM to convert are less than 2 TiB.
- Citrix Hypervisor pool (or host) has networks that the converted VMs use.

To convert your VMware ESXi/vCenter VMs into VMs that can run in a Citrix Hypervisor environment:

1. Ensure that the virtual appliance is installed and running on the Citrix Hypervisor server or pool where you want to import the VMs.
2. In XenCenter, go to **Pool > Conversion Manager**.

The **Conversion Manager** window opens. Wait while the wizard connects to your virtual appliance.

3. Click **New Conversion**.

4. 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 Citrix Hypervisor.
- **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.

5. In the **Virtual Machines** page, select from the list of VMs hosted in the VMware server the VMs that you want to convert. Click **Next**.

6. In 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.

7. On the **Networking** page, for each VMware network listed, select the Citrix Hypervisor network to map it to. You can also select whether to preserve virtual MAC addresses. Click **Next**.

8. Review the options you configured for the conversion process. You can click **Previous** to change these options. To proceed with the configuration shown, click **Finish**.

The conversion process begins. Conversion from ESXi or vSphere can take several minutes depending on the size of the virtual disks.

After converting your VMs, the Conversion Manager automatically shuts down by itself, saving resources on the host. Start a VM by selecting the VM's host and then clicking **Pool > Conversion Manager**.

The **Conversion Manager** window displays conversions in progress and completed conversions.

Steps after conversion

For Windows VMs, you must install XenServer VM Tools (formerly Citrix VM Tools) for Windows. For Linux VMs, you do not need to install XenServer VM Tools for Linux as the Conversion Manager automatically installs it during the conversion process.

In XenCenter, perform the following steps on your newly converted VMs:

On Windows machines:

1. 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.
2. On Windows VMs, install XenServer VM Tools (formerly Citrix VM Tools) for Windows to obtain high-speed I/O for enhanced disk and network performance. XenServer VM Tools for Windows also enable certain functions and features, including cleanly shutting down, rebooting, suspending, and live migrating VMs. You can download the XenServer VM Tools for Windows from the [Citrix Hypervisor downloads page](#).

If you are working with a VM that does not have XenServer VM Tools installed, a XenServer VM Tools not installed message appears on the **General** tab in the **General** pane.

Note:

XenServer VM Tools for Windows must be installed on each Windows VM for the VM to have a fully supported configuration. Although Windows VMs function without XenServer VM Tools for Windows, their performance can be impacted.

Enable VNC On Linux machines

On Linux VMs, configure the VNC server. For more information, see [Enable VNC for Linux VMs](#).

Note:

The VNC password must have at least six characters.

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.

Troubleshoot XenServer Conversion Manager

March 6, 2024

This section provides information about troubleshooting the conversion process and converted VMs.

Problems starting a converted VM

In general, conversion runs smoothly and XenServer Conversion Manager (formerly Citrix Hypervisor 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 XenServer Conversion Manager virtual appliance was unable to configure a Windows device that is critical to boot in Citrix Hypervisor for the first time. Save the logs and send them to Citrix 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 Citrix Hypervisor, 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 Citrix Hypervisor. This issue occurs because the migration process assigns the IDE hard disks lower device numbers than SCSI disks. However, Citrix Hypervisor 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 XenServer Conversion Manager virtual appliance and retry.

Logs of failed conversions are stored in the XenServer Conversion Manager virtual appliance and can be retrieved by clicking **Fetch All Logs** on the **Conversion Manager** window. When you contact Citrix support to raise any issues, we recommend that you provide the conversion log file and, also, a full server status report for troubleshooting. For more information, see [Creating a Server Status Report](#).

Command-line interface

May 28, 2024

The xe CLI enables you to script and automate system administration tasks. Use the CLI to integrate Citrix Hypervisor into an existing IT infrastructure.

Getting started with the xe CLI

The xe command line interface is installed by default on all Citrix Hypervisor servers. A remote Windows version is included with XenCenter. A stand-alone remote CLI is also available for Linux.

On your Citrix Hypervisor server

The xe command line interface is installed by default on your server. 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 server where you want to run the command.
- SSH into the server 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)\Citrix\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 Citrix Hypervisor 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 Citrix Hypervisor server, 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 Citrix Hypervisor server.

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 Citrix Hypervisor 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 Citrix Hypervisor server (defaults to 443).

Example: On the local Citrix Hypervisor server:

```
1 xe vm-list
```

Example: On a remote Citrix Hypervisor server:

```
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 Citrix Hypervisor server:

```
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 Citrix Hypervisor server, 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 Citrix Hypervisor server 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 Citrix Hypervisor server, 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 Citrix Hypervisor server, 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 Citrix Hypervisor server. 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*
- *pbid*
- *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

Citrix Hypervisor 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 Citrix Hypervisor host from a remote instance of the `xe` CLI.

To create a secret object, run the following command on your Citrix Hypervisor host.

```
1 xe secret-create value=my-password
```

A secret is created and stored on the Citrix Hypervisor 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 Citrix Hypervisor 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
<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 master 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 Citrix Hypervisor server
- 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 Citrix Hypervisor servers.

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
<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

Parameter Name	Description	Type
<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 Citrix Hypervisor server 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.

Cluster commands

Commands for working with clustered pools.

Clustered pools are resource pools that have the clustering feature enabled. Use these pools with GFS2 SRs. For more information, see [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 clustered pools.

Cluster parameters

Clusters have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the cluster	Read only
<code>cluster-hosts</code>	A list of unique identifiers/object references for the hosts in the cluster	Read only set parameter
<code>cluster-token</code>	The secret key used by <code>xapi-clusterd</code> when it talks to itself on other hosts	Read only
<code>cluster-stack</code>	The technology stack providing the clustering capabilities. Possible values are corosync .	Read only
<code>allowed-operations</code>	Lists the operations allowed in this state. This list is advisory only and the 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 cluster state may have changed by the time a client reads this field.	Read only set parameter

Parameter Name	Description	Type
<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>	True if automatically joining new pool members to the cluster. This is set to true .	Read only
<code>cluster-config</code>	A list of key/value pairs that specify extra configuration parameters for the cluster.	Read only map parameter
<code>other-config</code>	A list of key/value pairs that specify extra configuration parameters for the cluster.	Read/write map parameter

cluster-host-destroy

```
1 xe cluster-host-destroy uuid=host_uuid
```

Destroy a cluster host, effectively leaving the cluster.

cluster-host-disable

```
1 xe cluster-host-disable uuid=cluster_uuid
```

Disable cluster membership for an enabled cluster host.

cluster-host-enable

```
1 xe cluster-host-enable uuid=cluster_uuid
```

Enable cluster membership for a disabled cluster host.

cluster-host-force-destroy

```
1 xe cluster-host-force-destroy uuid=cluster_host
```

Destroy a cluster host object forcefully, effectively leaving the 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 cluster.

cluster-pool-destroy

```
1 xe cluster-pool-destroy cluster-uuid=cluster_uuid
```

Destroy pool-wide cluster. The pool continues to exist, but it is no longer 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 cluster.

cluster-pool-resync

```
1 xe cluster-pool-resync cluster-uuid=cluster_uuid
```

Resync a 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 Citrix Hypervisor.

diagnostic-compact

```
1 xe diagnostic-compact
```

Perform a major GC collection and heap compaction.

DEPRECATED: diagnostic-db-log

```
1 xe diagnostic-db-log
```

Start logging the database operations. Warning: once started, this cannot be stopped.

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
pool	A pool of physical hosts
vm	A Virtual Machine
host	A physical host
network	A virtual network
vif	A virtual network interface
pif	A physical network interface (separate VLANs are represented as several PIFs)
sr	A storage repository
vdi	A virtual disk image

Class name	Description
<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
<code>gpu-group-uuid</code>	The unique identifier/object reference for the GPU group that this pGPU has been automatically assigned to by Citrix Hypervisor. 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 Citrix Hypervisor server to which the pGPU is connected	Read only
<code>host-name-label</code>	The name of the Citrix Hypervisor server 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

Parameter Name	Description	Type
<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.

gpu-group-param-get-uuid

```
1 xe gpu-group-param-get-uuid uuid=uuid_of_group param-name=supported-vGPU-types|enabled-vGPU-types
```

Returns the supported or enabled types for this GPU Group.

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
<code>currently-attached</code>	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
<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_vm param-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 Citrix Hypervisor server.

Citrix Hypervisor servers are the physical servers running Citrix Hypervisor 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 Citrix Hypervisor server 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 Citrix Hypervisor servers on which to perform the operation. The simplest is by supplying the argument `host=uuid_or_name_label`. You can also specify Citrix Hypervisor by filtering the full list of hosts on the values of fields. For example, specifying `enabled=true` selects all Citrix Hypervisor servers whose `enabled` field is equal to `true`. Where multiple Citrix Hypervisor servers match and the operation can be performed on multiple Citrix Hypervisor servers, 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 Citrix Hypervisor servers are given, the operation is performed on all Citrix Hypervisor servers.

Host parameters

Citrix Hypervisor servers have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the Citrix Hypervisor server	Read only
<code>name-label</code>	The name of the Citrix Hypervisor server	Read/write
<code>name-description</code>	The description string of the Citrix Hypervisor server	Read only

Parameter Name	Description	Type
<code>enabled</code>	Value is false if disabled. This prevents any new VMs from starting on the hosts and prepares the hosts to be shut down or rebooted. Value is true 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
<code>software-version</code>	List of versioning parameters and their values	Read only map parameter
<code>capabilities</code>	List of Xen versions that the Citrix Hypervisor server can run	Read only set parameter
<code>other-config</code>	A list of key/value pairs that specify extra configuration parameters for the Citrix Hypervisor server	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>	Citrix Hypervisor server host name	Read only
<code>address</code>	Citrix Hypervisor server IP address	Read only

Parameter Name	Description	Type
<code>license-server</code>	A list of key/value pairs that specify information about the license server. The default port for communications with Citrix products is 27000. For information on changing port numbers due to conflicts, see Change port numbers	Read only map parameter
<code>supported-bootloaders</code>	List of bootloaders that the Citrix Hypervisor server supports, for example, pygrub , eliloader	Read only set parameter
<code>memory-total</code>	Total amount of physical RAM on the Citrix Hypervisor server, 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>	True if the host is operational	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 server 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 server 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

Parameter Name	Description	Type
<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
<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

Citrix Hypervisor servers contain some other objects that also have parameter lists.

CPUs on Citrix Hypervisor servers 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 Citrix Hypervisor server	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

Parameter Name	Description	Type
<code>host-uuid</code>	The UUID of the host the CPU is in	Read only
<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 Citrix Hypervisor servers have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the crashdump	Read only
<code>host</code>	Citrix Hypervisor server 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=xenserver_edition=  
    "free" "per-socket" "xendesktop"]
```

Assigns the Citrix Hypervisor license to a host server. When you assign a license, Citrix Hypervisor 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 Citrix Hypervisor for Citrix Virtual Desktops and DaaS editions, use `"xendesktop"`.

For initial licensing configuration, see also `license-server-address` and `license-server-port`.

host-backup

```
1 xe host-backup file-name=backup_filename host=host_name
```

Download a backup of the control domain of the specified Citrix Hypervisor server 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 Citrix Hypervisor server.

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 Citrix Hypervisor servers, which prevents any new VMs from starting on them. This action prepares the Citrix Hypervisor servers 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 Citrix Hypervisor servers.

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 Citrix Hypervisor server. Use this command only if the Citrix Hypervisor server is in emergency mode. Emergency mode means that the host is a member in a resource pool whose master 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 Citrix Hypervisor server where the command is run.

host-enable

```
1 xe host-enable [host-selector=host_selector_value...]
```

Enables the specified Citrix Hypervisor servers, 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 master, then another host must be selected to be the pool master. To change the pool master 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 server, which causes HA to elect a new master 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 Citrix Hypervisor server 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 Citrix Hypervisor server to “forget” is dead. However, if the Citrix Hypervisor server 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" \
11    min-time="-1" pii="maybe"/>
12   ...
13 </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 servers even with no master server present.

host-license-add

```
1 xe host-license-add [license-file=path/license_filename] [host-uuid=
  host_uuid]
```

For Citrix Hypervisor (free edition), use to parse a local license file and add it to the specified Citrix Hypervisor server.

host-license-remove

```
1 xe host-license-remove [host-uuid=host_uuid]
```

Remove any licensing applied to a host.

host-license-view

```
1 xe host-license-view [host-uuid=host_uuid]
```

Displays the contents of the Citrix Hypervisor server 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 Citrix Hypervisor servers. 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 par-

tion 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 Citrix Hypervisor server the CLI is connected to. The command is not forwarded to the pool master when applied to a member Citrix Hypervisor server.

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 Citrix Hypervisor server 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 Citrix Hypervisor server immediately rebinds. This command works both in normal and emergency mode.

If the UUID of a PIF object is specified, the Citrix Hypervisor server 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 Citrix Hypervisor server.

host-power-on

```
1 xe host-power-on [host=host_uuid]
```

Turns on power on Citrix Hypervisor servers 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 Citrix Hypervisor servers. 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 Citrix Hypervisor servers are members of a pool, the loss of connectivity on shutdown is handled and the pool recovers when the Citrix Hypervisor servers returns. The other members and the master continue to function.

If you shut down the master, the pool is out of action until one of the following actions occurs:

- You make one of the members into the master
- The original master is rebooted and back on line.

When the master is back online, the members reconnect and synchronize with the master.

host-restore

```
1 xe host-restore [file-name=backup_filename] [host-selector=  
host_selector_value...]
```

Restore a backup named `file-name` of the Citrix Hypervisor server 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 Citrix Hypervisor server.

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 Citrix Hypervisor server 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 Citrix Hypervisor server. 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" | "DRAC" | "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 Citrix Hypervisor 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 Citrix Hypervisor servers. The specified Citrix Hypervisor servers 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 Citrix Hypervisor servers are members of a pool, the loss of connectivity on shutdown is handled and the pool recovers when the Citrix Hypervisor servers returns. The other members and the master continue to function.

If you shut down the master, the pool is out of action until one of the following actions occurs:

- You make one of the members into the master
- The original master is rebooted and back on line.

When the master is back online, the members reconnect and synchronize with the master.

If HA is enabled for the pool, one of the members is made into a master automatically. If HA is disabled, you must manually designate the desired server as master 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 master 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 Citrix Hypervisor servers. 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.

DEPRECATED: Log commands

Commands for working with logs.

DEPRECATED: log-get

```
1 xe log-get
```

Return the log currently stored in the string logger.

DEPRECATED: log-get-keys

```
1 xe log-get-keys
```

List the keys known by the logger.

DEPRECATED: log-reopen

```
1 xe log-reopen
```

Reopen all loggers (use this for rotating files).

DEPRECATED: log-set-output

```
1 xe log-set-output output=output [key=key] [level=level]
```

Set all loggers to the specified output (*nil*, *stderr*, *string*, *file:file name*, *syslog:something*).

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
<i>uuid</i>	The unique identifier/object reference for the message	Read only
<i>name</i>	The unique name of the message	Read only
<i>priority</i>	The message priority. Higher numbers indicate greater priority	Read only
class	The message class, for example VM.	Read only
<i>obj-uuid</i>	The uuid of the affected object.	Read only
<i>timestamp</i>	The time that the message was generated.	Read only
<i>body</i>	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 Citrix Hypervisor servers to this network	Read only set parameter

Parameter Name	Description	Type
<code>bridge</code>	Name of the bridge corresponding to this network on the local Citrix Hypervisor server	Read only
<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 Citrix Hypervisor server 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 no to disable autonegotiation of the physical interface or bridge. Default is yes.	Read/write
<code>other-config:ethtool-rx</code>	Set to on to enable receive checksum, off to disable	Read/write
<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

Parameter Name	Description	Type
<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>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
<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 server.

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 prechecks 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 server.

PBD commands

Commands for working with PBDs (Physical Block Devices). PBDs are the software objects through which the Citrix Hypervisor server 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>	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 Citrix Hypervisor server. 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 Citrix Hypervisor server. If this command succeeds, the referenced SR (and the VDIs contained within) becomes visible to the Citrix Hypervisor server.

pbd-unplug

```
1 xe pbd-unplug uuid=uuid_of_pbd
```

Attempt to unplug the PBD from the Citrix Hypervisor server.

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, eth0)	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

Parameter Name	Description	Type
<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
<code>MTU</code>	Maximum Transmission Unit of the PIF in bytes.	Read only
<code>VLAN</code>	VLAN tag for all traffic passing through this interface. -1 indicates that no VLAN tag is assigned	Read only
<code>bond-master-of</code>	The UUID of the bond this PIF is the master of (if any)	Read only
<code>bond-slave-of</code>	The UUID of the bond this PIF is part of (if any)	Read only
<code>management</code>	Is this PIF designated to be a management interface for the control domain	Read only
<code>network-uuid</code>	The unique identifier/object reference of the virtual network to which this PIF is connected	Read only
<code>network-name-label</code>	The name of the virtual network to which this PIF is connected	Read only
<code>host-uuid</code>	The unique identifier/object reference of the Citrix Hypervisor server to which this PIF is connected	Read only
<code>host-name-label</code>	The name of the Citrix Hypervisor server to which this PIF is connected	Read only
<code>IP-configuration-mode</code>	Type of network address configuration used; DHCP or static	Read only
<code>IP</code>	IP address of the PIF. Defined here when IP-configuration-mode is static; undefined when DHCP	Read only

Parameter Name	Description	Type
<code>netmask</code>	Netmask of the PIF. Defined here when IP-configuration-mode is static; undefined when supplied by DHCP	Read only
<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

Parameter Name	Description	Type
<code>other-config:ethtool-duplex</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
<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:bondndelay</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>	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. Alternatively, 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 Citrix Hypervisor server.

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 Citrix Hypervisor server.

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 Citrix Hypervisor servers. 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 Citrix Hypervisor server is really a pool consisting of a single member by default. When your Citrix Hypervisor server is joined to a pool, it is designated as a member, and the pool it has joined becomes the master for 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 Citrix Hypervisor server designated as the pool's master	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>	True if HA is enabled for the pool, false otherwise	Read only
<code>ha-configuration</code>	Reserved for future use.	Read only

Parameter Name	Description	Type
<code>ha-statefiles</code>	Lists the UUIDs of the VDIs being used by HA to determine storage health	Read only
<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>	True if the pool is allowed to be overcommitted, False otherwise	Read/write
<code>ha-overcommitted</code>	True if the pool is overcommitted	Read only
<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 Citrix Hypervisor 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

pool-apply-edition

```
1 xe pool-apply-edition edition=edition [uuid=uuid] [license-server-address=address] [license-server-port=port]
```

Apply an edition across the pool.

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 master 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 Citrix Hypervisor server to become the master of an existing pool. This command performs an orderly handover of the role of master host to another host in the resource pool. This command only works when the current master 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 Citrix Hypervisor server 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 server to reset its master server address to the new value and attempt to connect to it. Do not run this command on master servers.

pool-emergency-transition-to-master

```
1 xe pool-emergency-transition-to-master
```

Instruct a member Citrix Hypervisor server to become the pool master. The Citrix Hypervisor server accepts this command only after the host has transitioned to emergency mode. Emergency mode means it is a member of a pool whose master 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=wb_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 Citrix Hypervisor server 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 Citrix Hypervisor master host 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 Citrix Hypervisor server, 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 Citrix Hypervisor server.

pool-recover-slaves

```
1 xe pool-recover-slaves
```

Instruct the pool master to try to reset the master 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 master.

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 master 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 servers in a pool that enables the server to prove its membership to a pool. Users with the Pool Admin role can view this secret when connecting to the server 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.

igmp-snooping-enabled

```
1 xe pool-param-set [uuid=pool-uuid] [igmp-snooping-enabled=true|false]
```

Enables or disables IGMP snooping on a Citrix Hypervisor pool.

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 Citrix Hypervisor 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
<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

Parameter Name	Description	Type
<code>required-api-version</code>	Minimum SM API version required on the Citrix Hypervisor server	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

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
<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

Parameter Name	Description	Type
<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>	True if this SR can be shared between multiple Citrix Hypervisor servers. False otherwise.	Read/write
<code>introduced-by</code>	The <code>drtask</code> (if any) which introduced the SR	Read only
<code>is-tools-sr</code>	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>	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>	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 Citrix Hypervisor server. If `shared` is set to **true**, a PBD is created for each Citrix Hypervisor server in the pool. If `shared` is not specified or set to **false**, a PBD is created only for the Citrix Hypervisor server 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 Citrix Hypervisor server.

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 Citrix Hypervisor server. 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

Parameter Name	Description	Type
<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
<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

Parameter Name	Description	Type
<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. Citrix Hypervisor 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.

You can tune a vCPU's pinning with:

```
1 xe template-param-set uuid=<template_uuid> vCPUs-params:mask=1,2,3
```

A VM created from this template run on physical CPUs 1, 2, and 3 only.

You can also tune the vCPU priority (xen scheduling) with the cap and weight parameters. For example:

```
1 xe template-param-set uuid=<template_uuid> VCPUs-params:weight=512 xe template-param-set uuid=<template_uuid> VCPUs-params:cap=100
```

A VM based on this template with a weight of 512 get twice as much CPU as a domain with a weight of 256 on a contended host. Legal weights range from 1 to 65535 and the default is 256.

The cap optionally fixes the maximum amount of CPU a VM based on this template can consume, even if the Citrix Hypervisor server has idle CPU cycles. The cap is expressed in percentage of one physical CPU: 100 is 1 physical CPU, 50 is half a CPU, 400 is 4 CPUs, and so on. The default, 0, means that there is no upper cap.

- `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 HVM guests (for example, Windows guests):

```
1 xe vm-param-set uuid=<vm_uuid> platform:parallel=none
```

To disable the emulation of a serial port for HVM guests:

```
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 for HVM guests:

```
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 0–15. 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 0–15. This list is informational only, and other devices may be used (but may not work).
- `HVM-boot-policy` (read/write) the boot policy for HVM guests. Either BIOS Order or an empty string.

- `HVM-boot-params` (read/write map parameter) the order key controls the HVM 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 Citrix Hypervisor server on which a VM based on this template is resident. Appears as `not in database` for a template
- `affinity` (read/write) the Citrix Hypervisor server 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
- `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) 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) 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 Citrix Hypervisor server 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

Citrix Hypervisor server updates have the following parameters:

Parameter Name	Description	Type
<code>uuid</code>	The unique identifier/object reference for the update	Read only

Parameter Name	Description	Type
<code>host</code>	The list of hosts that this update is applied to	Read only
<code>host-uuid</code>	The unique identifier for the Citrix Hypervisor server 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 Citrix Hypervisor server. 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 prechecks contained within the specified update on the specified Citrix Hypervisor server.

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 Citrix Hypervisor servers 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>	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>	True if the VBD is attached on this host, false otherwise	Read only
<code>storage-lock</code>	True if a storage-level lock was acquired	Read only

Parameter Name	Description	Type
<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>	True if this VBD supports hot unplug	Read/write
<code>attachable</code>	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 0–15, 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
uuid	The unique identifier/object reference for the VDI	Read only
name-label	The name of the VDI	Read/write
name-description	The description string of the VDI	Read/write
allowed-operations	A list of the operations allowed in this state	Read only set parameter
current-operations	A list of the operations that are currently in progress on this VDI	Read only set parameter
sr-uuid	SR in which the VDI resides	Read only
vbd-uuids	A list of VBDs that refer to this VDI	Read only set parameter
crashdump-uuids	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>	True if this VDI may be shared	Read only
<code>read-only</code>	True if this VDI can only be mounted read-only	Read only
<code>storage-lock</code>	True if this VDI is locked at the storage level	Read only
<code>parent</code>	References the parent VDI when this VDI is part of a chain	Read only
<code>missing</code>	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>	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>/device-id/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
<code>is-a-snapshot</code>	True if this VDI is a VM storage snapshot	Read only
<code>snapshot_of</code>	The UUID of the storage this VDI is a snapshot of	Read only

Parameter Name	Description	Type
<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 Citrix Hypervisor 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) 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`, Citrix Hypervisor applies a filtering rule so that the VIF cannot send or receive traffic. If the default-lockingmode is set to `unlocked`, Citrix Hypervisor 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) Citrix Hypervisor 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 0-15.

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 Citrix Hypervisor server.

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 on a pool, by determining which interface (for example, `eth0`) the specified network is on (on each host) and 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) 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) 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.

You can tune a vCPU's pinning with

```
1 xe vm-param-set uuid=<vm_uuid> VCPUs-params:mask=1,2,3
```

The selected VM then runs on physical CPUs 1, 2, and 3 only.

You can also tune the vCPU priority (xen scheduling) with the cap and weight parameters. For example:

```
1 xe vm-param-set uuid=<vm_uuid> VCPUs-params:weight=512 xe vm-param-set uuid=<vm_uuid> VCPUs-params:cap=100
```

A VM with a weight of 512 get twice as much CPU as a domain with a weight of 256 on a contended Citrix Hypervisor server. Legal weights range from 1 to 65535 and the default is 256. The cap optionally fixes the maximum amount of CPU a VM will be able to consume, even if the Citrix Hypervisor server has idle CPU cycles. The cap is expressed in percentage of one physical CPU: 100 is 1 physical CPU, 50 is half a CPU, 400 is 4 CPUs, and so on The default, 0, means that there is no upper cap.

- `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 0–15. 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 0–15. This list is informational only, and other devices may be used (but might not work).
- `HVM-boot-policy` (read/write) the boot policy for HVM guests. Either BIOS Order or an empty string.
- `HVM-boot-params` (read/write map parameter) the order key controls the HVM 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 Citrix Hypervisor server on which a VM is resident
- `affinity` (read/write) The Citrix Hypervisor server which the VM has preference for running on. Used by the `xe vm-start` command to decide where to run the VM
- `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 Citrix Hypervisor server.
- `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) 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) 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) True if the VM is running. 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 the value of the `allowed-VBD-devices` parameter of the 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-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 Citrix Hypervisor server 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 [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 `host` parameter can be either the name or the UUID of the Citrix Hypervisor server. 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 master, 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:vif_uuid=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](#).

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 master 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-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 Citrix 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
name-label	Name of the snapshot schedule.	Read/write
name-description	Description of the snapshot schedule.	Read/write
type	Disk snapshot or memory snapshot.	Read/write
frequency	Hourly; Daily; Weekly	Read/write

Parameter Name	Description	Type
<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 pusb-param-set uuid=pusb_uuid passthrough-enabled=true/false
```

Enable/disable USB Pass-through.

pusb-scan

```
1 xe pusb-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.

Common criteria

January 9, 2023

Citrix Hypervisor 8.2 Cumulative Update 1 is Common Criteria certified EAL2+. For more information, see [Citrix Common Criteria Certification Information](#).

- [Citrix Hypervisor 8.2 Cumulative Update 1 LTSR Common Criteria Evaluated Configuration Guide \(PDF\)](#)
- [Citrix Hypervisor 8.2 Cumulative Update 1 LTSR Product Documentation \(PDF\)](#)
- [Citrix Hypervisor 8.2 Cumulative Update 1 LTSR Management API \(PDF\)](#)

Third party notices

November 22, 2022

This release of Citrix Hypervisor includes third-party software licensed under a number of different licenses.

To extract the licensing information from your installed Citrix Hypervisor product and components, see the instructions in [Citrix Hypervisor 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 (eay@cryptsoft.com).

- Citrix Hypervisor High Availability is powered by everRun, a registered trademark of Stratus Technologies Bermuda, Limited.

Citrix Hypervisor Open Source Licensing and Attribution

July 17, 2024

The Citrix Hypervisor 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 Citrix Hypervisor 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 Citrix Hypervisor server

This article provides a method to extract the licensing information from all RPM packages included in your Citrix Hypervisor installation.

Get overview information

To list all RPMs and their licenses:

1. Connect to your Citrix Hypervisor server 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 Citrix Hypervisor server 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   : 2.2.0
3 Release   : 1.xs8
4 Architecture: noarch
5 Install Date: Thu 03 Jun 2021 08:36:59 AM UTC
6 Group      : Unspecified
7 Size       : 97131
8 License    : GPL
9 Signature  : (none)
10 Source RPM : host-upgrade-plugin-2.2.0-1.xs8.src.rpm
11 Build Date : Fri 09 Oct 2020 02:58:51 PM UTC
12 Build Host : 2da9e81a970c4f02af07e64918d7f5f3
13 Relocations : (not relocatable)
14 Packager   : Koji
15 Vendor     : Citrix Systems
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: Thu 03 Jun 2021 08:36:22 AM UTC
25 Group      : Applications/Text
26 Size       : 525707
27 License    : GPLv3+
28 Signature  : RSA/SHA256, Wed 25 Nov 2015 03:16:04 PM UTC, Key ID
                24c6a8a7f4a80eb5
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.
```

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 Citrix Hypervisor product. To view the license text for these components, you can look inside the source RPMs. Citrix makes the source RPMs for the Citrix Hypervisor server available in the following locations

- For the initial product release, source files are provided on the [product download page](#).
- For any updates or hotfixes to the initial release, updated source files are provided in the corresponding article on the [Citrix Support site](#). (Go to **Common Issues > XenServer Software Updates**.)

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
```

Multiple licenses

Some components in the Citrix Hypervisor 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.

Other Citrix Hypervisor components

Supplemental Packs

Supplemental packs are installed into the Citrix Hypervisor server. If you have supplemental packs installed in your server, 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 [product download 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**.

The XenCenter source files are also provided on the [product download page](#).

XenServer VM Tools for Windows

The XenServer VM Tools for Windows (formerly Citrix VM Tools) 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\Citrix\XenTools\Drivers***.inf`.

Source is not provided for XenServer VM Tools for Windows.

Citrix VM Tools for Linux

The Citrix 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 Citrix Hypervisor environment:

- Demo Linux Virtual Appliance
- XenServer Conversion Manager (formerly Citrix Hypervisor 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 Citrix Hypervisor server 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 XenServer Conversion Manager virtual appliance and Workload Balancing virtual appliance dynamically use some third-party components.

- For XenServer Conversion Manager virtual appliance, the license files for these components are located at the following path: `/opt/vpxxcm/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 [Citrix Hypervisor product downloads page](#).

Develop for XenServer

June 20, 2024

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](#) or [Citrix Hypervisor 8.2 Cumulative Update 1](#).

Other development tools

To develop an incremental backup solution, see [Changed Block Tracking](#).

To develop a driver disk or supplemental pack, see [Supplemental Pack and DDK Guide](#).

To develop a plugin for XenCenter, see [XenCenter Plugin Specification](#).

Data Governance

August 31, 2023

This article provides information regarding the collection, storage, and retention of logs by Citrix Hypervisor.

Citrix Hypervisor is a server virtualization platform that enables the customer to create and manage a deployment of virtual machines. XenCenter is the management UI for Citrix Hypervisor. Citrix Hypervisor and XenCenter can collect and store customer data as part of providing the following capabilities:

- **Server status reports** - A server status report can be generated on-demand and uploaded to Citrix Insight Services or provided to Citrix Support. The server status report contains information that can aid in diagnosing issues in the customer's environment.
- **Automatic updates for the Management Agent** - The Management Agent runs within VMs hosted on a Citrix Hypervisor server 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 check for updates** - This feature determines whether any hotfixes, cumulative updates, or new releases are available for the Citrix Hypervisor servers and pools XenCenter manages. As part of checking for updates, this feature makes a web request to Citrix that includes telemetry. This telemetry is not user-specific and is used to estimate the total number of XenCenter instances worldwide.
- **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.

Any information received by Cloud Software Group is treated in accordance with our [Agreements](#).

Server status reports

During the course of operation a Citrix Hypervisor server collects and logs various information on the server where Citrix Hypervisor is installed. These logs can be collected as part of a server status report.

A server status report can be generated on-demand and uploaded to Citrix Insight Services or provided to Citrix Support. The server status report contains information that can aid in diagnosing issues in the customer's environment.

Server status reports that are uploaded to Citrix Insight Services are stored in Amazon S3 environments located in the United States.

Citrix Hypervisor and XenCenter collect information from the following data sources:

- XenCenter
- Citrix Hypervisor servers 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>xapi-debug</code>	maybe
<code>xen-info</code>	maybe
<code>conntest</code>	no
<code>xha-liveset</code>	maybe
<code>high-availability</code>	maybe
<code>firstboot</code>	yes
<code>xenserver-databases</code>	yes
<code>multipath</code>	maybe
<code>disk-info</code>	maybe
<code>xenserver-logs</code>	maybe
<code>xenserver-install</code>	maybe
<code>process-list</code>	yes
<code>blobs</code>	no
<code>xapi</code>	yes
<code>host-crashdump-logs</code>	maybe
<code>xapi-subprocess</code>	no
<code>pam</code>	no
<code>control-slice</code>	maybe
<code>tapdisk-logs</code>	no
<code>kernel-info</code>	maybe
<code>xenserver-config</code>	maybe
<code>xenserver-domains</code>	no
<code>device-model</code>	yes

Log type	Contains PII?
hardware-info	maybe
xenopsd	maybe
loopback-devices	maybe
system-services	no
system-logs	maybe
network-status	yes
v6d	maybe
CVSM	no
message-switch	maybe
VM-snapshot-schedule	no
xcp-rrdd-plugins	maybe
yum	if customized
fcoe	yes
xapi-clusterd	maybe
network-config	if customized
boot-loader	no

Management Agent automatic updates

The Management Agent runs within VMs hosted on a Citrix Hypervisor server 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.

The web logs captured from the requests made by the Management Agent automatic updates feature are located in a Microsoft Azure Cloud environment located 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	
VM UUID	A unique user ID for the VM where the Management Agent is installed	

XenCenter check for updates

This feature determines whether any hotfixes, cumulative updates, or new releases are available for the Citrix Hypervisor servers and pools XenCenter manages. As part of checking for updates, this feature makes a web request to Cloud Software Group that includes telemetry. This telemetry does not personally identify users and is used to estimate the total number of XenCenter instances world-wide.

The web logs captured from the requests made by the XenCenter check for updates feature are located in a Microsoft Azure Cloud environment located in the United States. These logs are then copied to a log management server in the United Kingdom.

The web requests made by the XenCenter check for updates feature are made over HTTPS. Web log files are transmitted securely to the log management server.

The XenCenter check for updates feature is enabled by default. You can choose to disable this feature.

Data collected

The check for updates feature web requests contain the following data points:

Data collected	Description	What we use it for
IP address	The IP address of the XenCenter host machine	
XenCenter version	The version of XenCenter making the request	

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



© 2024 Cloud Software Group, Inc. All rights reserved. Cloud Software Group, the Cloud Software Group logo, and other marks appearing herein are property of Cloud Software Group, Inc. and/or one or more of its subsidiaries, and may be registered with the U.S. Patent and Trademark Office and in other countries. All other marks are the property of their respective owner(s).