



# XenServer TechZone

## **Contents**

<b>TechZone</b>	<b>2</b>
<b>Reference architecture for Citrix workloads</b>	<b>2</b>

## TechZone

July 1, 2024

Welcome to the XenServer TechZone!

XenServer Tech Zone is home to technical, in-depth articles inspired and driven by technical communities and enthusiasts. Whether you are an architect, consultant, engineer, or technical IT manager, you have come to the right place! Here, you will learn from experts and gain insights into XenServer and Citrix technologies and integrations for successful deployments with your Citrix environment through articles, tech best practices, videos, and much more.

## Reference architecture for Citrix workloads

July 1, 2024

This document serves as a blueprint for deploying XenServer to run Citrix workloads for the most common commercial-sized deployment that can scale from a few hundred to a few thousand VDAs. This reference architecture is valid whether using Citrix Virtual Apps and Desktops or Citrix DaaS. Enterprise-sized deployments may have additional considerations not covered in this reference architecture. Use [XenServer product documentation](#) along-side this document.

### Blueprint

#### Host and resource pool layer

- XenServer hosts should be part of a [resource pool](#) with a recommended maximum of 16 hosts.
- XenServer hosts in the same resource pool must have the same vendor, model, and features as the CPUs, as well as the same amount of memory.
- Memory should not be overcommitted. You need as much memory in a host as the VMs have allocated.
- See the Citrix Provisioning workloads section for requirements for local storage requirements as well as host memory considerations.

#### Network layer

- XenServer hosts should have network card speeds of 10 Gbps or greater.

- XenServer hosts should have a minimum of 4 network cards: 2 bonded pairs with 1 pair dedicated to storage traffic, and 1 pair used for the VM and management traffic.
- VLANs on external switches can be used to provide additional separation of storage, VM, and management traffic to meet security best practices, if desired.

### Storage layer

- Shared storage is recommended to ensure VMs can be migrated between hosts.
- NFS or SMB is recommended when using Machine Creation Service (MCS).
- Any [supported storage option](#) works when using Citrix Provisioning.
- Isolate storage networking traffic as outlined in the Network Layer section.

### Citrix Image Provisioning layer

Citrix Machine Creation Services (MCS) and Citrix Provisioning Services can be used separately or in combination for provisioning VDAs to XenServer.

**Citrix Provisioning workloads** If using Citrix Provisioning, we recommend that you enable the XenServer feature [PVS-Accelerator](#).

- 5 GB of cache on each host is recommended per vDisk version that you actively use.
- Memory cache is recommended instead of disk cache, if enough memory is available.
  - If using disk cache, local storage is recommended.

**Machine Creation Services workloads** If using Citrix Machine Creation Services, we recommend that you use both [Intellicache](#) and [storage read caching](#).

#### Intellicache:

- Enable Intellicache during the installation of XenServer by selecting “Enable thin provisioning (Optimized storage for Virtual Desktops)”.
- Intellicache uses local storage for the cache.
- XenServer hosts should have enterprise grade SSDs or NVME drives that support 512 byte sectors (or can emulate 512 byte sectors).
- NFS or SMB shared storage is recommended as it is required for the VDAs to allow a fully thin-provisioned solution with IntelliCache.
- When creating the Hosting Connection from Citrix, make sure the IntelliCache option is selected.

#### Storage read caching:

- On each XenServer host, increase Dom0 memory by 10 GB to provide space.

## Design decisions

This section provides more details about the reasons for the blueprint configuration as well as other potential configuration options.

### Host and resource pool layer

- While XenServer can support a resource pool with up to 64 hosts, limiting the resource pool to 16 hosts ensures that the time needed to perform updates (even when host reboots are required) is achievable within a working day. Additionally, there is an increased resilience to failures and the impact of failure (should one occur) is constrained to this set of hosts.
- When allocating VMs to a XenServer resource pools, ensure that there is enough capacity to operate all VMs with at least 1 host unavailable. This allows for maintenance operation on the pools without requiring VM outages.
- If XenServer hosts in the same Resource Pool have different amounts of memory, the XenServer host with the smallest amount of memory must be able to support workloads that get placed on it during failover scenarios or upgrades.
- XenServer hosts within the same resource pool should be on the same network, in the same datacenter or physical location, and only separated by a L2 switch (not a router).
- Create a separate resource pool for each set of XenServer hosts that are on a different network or in a different physical location.
- XenServer HA is not recommended for Citrix workloads/VDAs.
  - Protection at the VM level is not typically required due to the inherent manner in which Citrix Virtual Apps and Desktops workloads are dynamically created and destroyed
  - HA in a Citrix Virtual Apps and Desktops deployment can be beneficial to handle hardware failure or hypervisor crash. However when HA is enabled there is an increased risk of any temporary interruptions (in network or storage infrastructure) causing a host to ‘fence’ for safety resulting in an interruption of services (for end users) that otherwise might not have occurred.
- If possible, splitting VDAs over multiple pools ensures for availability in case of a pool failure.
- The total number of vCPUs allocated to VMs on any one host should not exceed the number of physical CPU threads for the host.

### Network layer

Other network card options for your hosts:

- 6 Network cards: 3 bonded pairs with 1 pair dedicated to storage traffic, 1 pair dedicated to VM traffic, and 1 pair dedicated to management traffic.
- 3 network cards: 1 network card dedicated to storage traffic, 1 network card dedicated to VM traffic, and 1 network card dedicated to management traffic.
- 2 network cards: 1 network card dedicated to the storage traffic, and 1 network card used for the VM and management traffic.

### **Citrix Provisioning Layer**

- Minimize the number of different golden images used in each resource pool to make best use of the caching technologies available.

Each image makes use of the caches. The more golden images there are the more likely the caches are to become full and become less effective. Making the caches bigger may also help with this aspect as the number of golden images increases.

### **Intellicache**

With Intellicache, if you are using block based storage, we recommend that you use full provision (LVM) mode. This mode is compatible with IntelliCache (which will enable faster VM operation with older/slower storage devices). Some block storage filers provide thin provisioning which can be used but care must be taken to avoid out-of-space conditions.

### **Reference materials**

- XenServer Product Documentation: <https://docs.xenserver.com/en-us/xenserver/8/>
- XenServer Technical Overview: <https://docs.xenserver.com/en-us/xenserver/8/technical-overview>
- Getting Started with XenCenter: <https://docs.xenserver.com/en-us/xencenter/current-release/intro-welcome>



© 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).