

TrueNAS® for VMware®

Version 1.2
Published 2011, updated 2021



Storage and Servers Driven by Open Source



Contents

| | |
|---|---|
| 1 Executive Summary | 1 |
| 2 History of iXsystems | 1 |
| 2.1 Open Source Background | 1 |
| 3 vSphere | 2 |
| 3.1 VMware HA | 3 |
| 3.2 Virtual Desktops | 3 |
| 3.2.1 VMware Horizon for VDI | 3 |
| 3.2.2 Citrix XenDesktop for VDI | 4 |
| 3.3 Server/Application VMs | 5 |
| 3.4 VMotion | 6 |
| 3.5 VAAI | 6 |
| 3.6 Best Practices for VAAI | 7 |
| 4 TrueNAS® Overview and Portfolio | 8 |
| 5 Conclusion | 8 |

©2011-2021 iXsystems, Inc. All rights reserved. TrueNAS, TrueOS, FreeNAS, TrueFlash and all associated logos and designs are trademarks or registered trademarks of iXsystems Corporation. All other registered trademarks or trademarks are property of their respective owners. This White Paper is for informational purposes only. **IXSYSTEMS MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS WHITE PAPER.**

1 Executive Summary

Hypervisors are used to deploy servers and desktops inside virtual machines (VMs) that are hosted and managed from a central location utilizing shared and VM-optimized storage. This provides cost-effective high availability and scalable performance. TrueNAS adds additional value to a virtualized environment by protecting it from silent data corruption.

TrueNAS is a storage array that uses the OpenZFS file system. OpenZFS was specifically designed to ensure the integrity of data. It uses features like end-to-end checksumming and copy-on-write to protect data against the silent data corruption caused by bit rot, current spikes, bugs in disk firmware, phantom writes, misdirected reads/writes, DMA parity errors, driver errors, disk errors, and accidental overwrites. No other file system, volume manager, or hardware RAID solution provides sufficient protection against such problems. OpenZFS also uses advanced read and write caching algorithms, allowing applications to get the performance of an all-flash array with the capacity of a disk array.

Enterprise IT infrastructures are powered by storage and data is arguably the most precious commodity companies possess. TrueNAS is VMware Ready certified and offers unified and scalable storage solutions to protect data, simplify its management, reduce operational costs, and optimize the performance of a virtualized environment. It unifies SAN and NAS in one array and provides integration with VMWare VAAI, OpenZFS snapshots, and thin provisioning. Easy to deploy and maintain, TrueNAS helps to ensure that critical business data remains available. Through innovative double-fault protection and the use of OpenZFS, storage is protected from disk, network, and storage node faults.

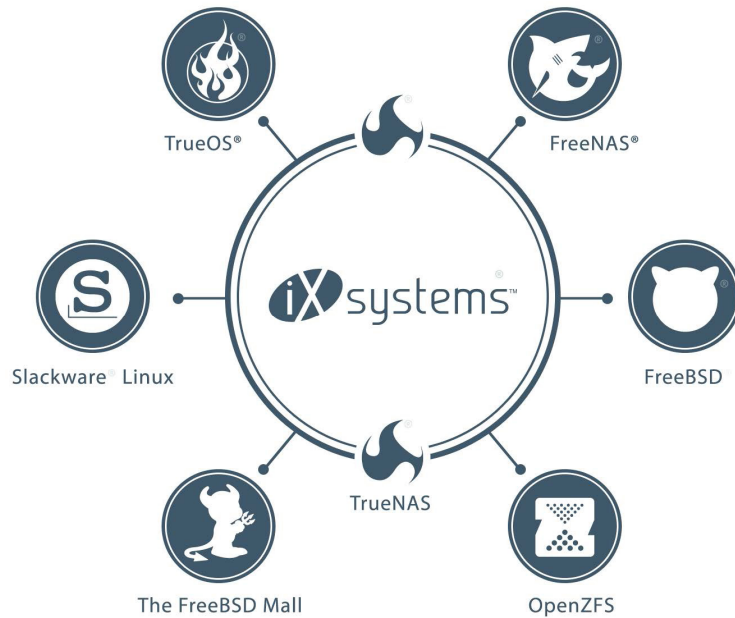
When coupled with VMware, TrueNAS coordinates the conversion of VMware snapshots to OpenZFS snapshots to create stable local or remote resurrection points. With its dynamic caching and storage optimization technology, TrueNAS eliminates the need to rely on multiple spindles of power hungry or short-stroked disks to get the IOPS needed for a VM deployment. Variable block sizes, high performance storage optimization, snapshots, replication, and flash I/O acceleration all increase the value provided by TrueNAS. Deploying a TrueNAS Enterprise storage array reduces overall cost and management overhead per VM.

2 History of iXsystems

Located in the heart of Silicon Valley, iXsystems is committed to serving technology needs with a focus on Open Source and enterprise hardware since our beginning in 1996. From our inception onward, we've been perfecting our craft of making quality storage solutions and custom-built servers backed by a passion for superior customer experience. See what our customers say about us at <https://www.vendop.com/vendor/ixsystems-inc/reviews/>.

2.1 Open Source Background

If your company leverages Open Source technology, it's a distinct advantage to work with a company that speaks your language. Nearly everything we do at iXsystems involves and benefits Open Source technology. We incorporate Open Source solutions into our storage and server product lines, and use them extensively ourselves. We are the main developers of TrueNAS® CORE, TrueNAS Enterprise, and TrueCommand® and contribute to FreeBSD, OpenZFS, and Slackware Linux. We spread the Open Source message in print, radio, and video publications and through our participation in various industry events around the globe.



We employ a long list of FreeBSD and Open Source project committers who came to iXsystems because of their passion for Open Source. This allows us to be successful, which in turn helps our clients do the same.

TrueNAS Enterprise is based on TrueNAS CORE, an Open Source Network Attached Storage software also developed by iXsystems. TrueNAS is a unified operating system designed specifically to provide NAS/SAN storage features. No additional software is required to make it fully functional. TrueNAS is based on the FreeBSD operating system and uses the OpenZFS file system as the basis for many its most powerful features.

TrueNAS is developed with extensive feedback and bug testing from the public. TrueNAS is the world's most popular Open Source storage project with millions of downloads and deployments in various environments around the world. TrueNAS CORE allows us to improve quality and ensure stable TrueNAS Enterprise products.

3 vSphere

VMware vSphere provides a suite of tools for managing ESX/ESXi-based cloud computing. TrueNAS supports vSphere 4.1 and later.

Gartner¹ states that 70% of enterprise customers have or will implement virtualization solutions. Of these, almost 80% are using VMware. These customers have transformed their physical servers to virtual servers due to density, management, control, security, or cost reasons.

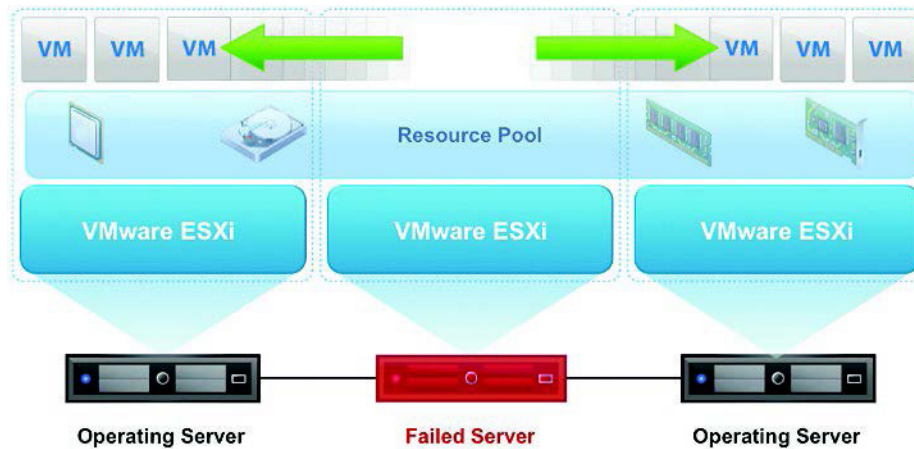
Log.in Consultants in their Project Virtual Reality Check² surveyed over 500 enterprise customers and found that customers are virtualizing their server applications and over 80% are also virtualizing Server-Based Computing (SBC) deployments such as XenApp. Over half of these customers virtualize using vSphere, over one third use Citrix Xen-Server, and a little more than ten percent use Hyper-V to deploy virtual desktops.

¹ <https://www.gartner.com/doc/3400418/magic-quadrant-x-server-virtualization>

² <https://www.projectvrc.com/white-papers/17-state-of-the-vdi-and-sbc-union-2015-v1.pdf>

3.1 VMware HA

VMware High Availability (HA) provides high availability for VMs and servers. In the event of physical server failure, affected virtual machines are automatically restarted on other production servers with spare capacity. In the case of a VM failure, VMware HA restarts the affected virtual machine on the same physical server.



3.2 Virtual Desktops

Most enterprises struggle to keep up with the proliferation and management of personal computers in their environment. Each computer must be managed as an individual entity. The benefits of distributed processing come at the cost of increased management. It costs time and money to set up, update, support, and ultimately decommission each computer. The initial cost of the machine is often dwarfed by operational costs. Having virtualized their server applications, many enterprises are now deploying SBC using Citrix XenApp or Microsoft RDS or are deploying virtual desktops using VMware Horizon or Citrix XenDesktop.

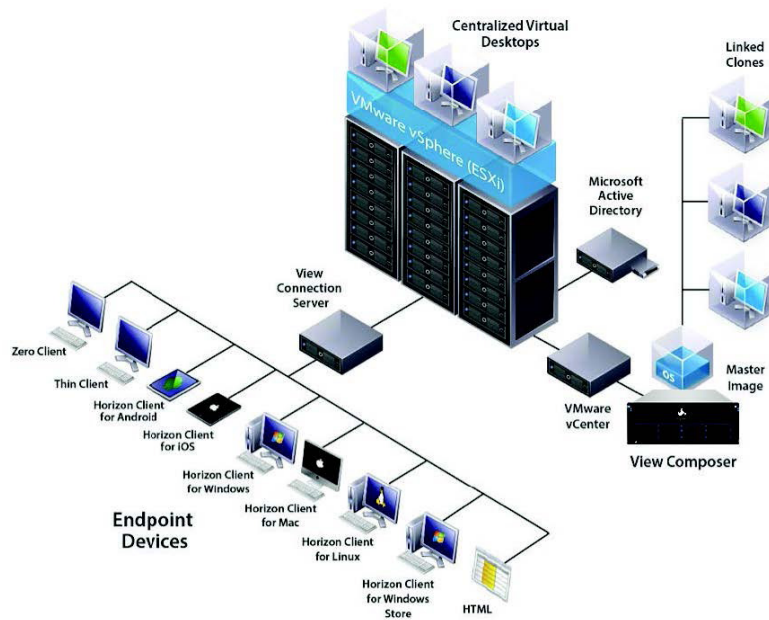
3.2.1 VMware Horizon for VDI

VMware Horizon provides remote desktop capabilities to users using VMware ESXi. A client desktop operating system such as Windows, Linux, or BSD runs within a virtual environment on a server.

When using VMware Horizon, iXsystems recommends the following best practices for optimizing data integrity and performance:

- Use TrueNAS 11.3 or later, which supports VAAI. VAAI is VMware's API framework that enables certain storage tasks, such as large data moves, to be offloaded from the virtualization hardware to the storage array.
- VMware Horizon relies on a SQL Server running on a server in the vSphere cluster. Configure this server to use TrueNAS to store its logs and databases in order to leverage the data integrity capabilities of OpenZFS.
- Enable desktop and email searching. While many Horizon Best Practices documents recommend disabling search for performance reasons, TrueNAS read caching ensures that search is nearly instantaneous.
- Provision the data pools and SQL Server storage using iSCSI. Define sync=always on each iSCSI dataset to ensure that every database transaction is written and flushed to stable storage.
- Use the TrueNAS VMware-Snapshot utility to protect VDI VMs. TrueNAS will automatically snapshot all running VMs before taking an OpenZFS snapshot of the storage backing that VMware datastore.
- Review the VMware Horizon Best Practices technical white paper³.

³ <http://www.vmware.com/files/pdf/view/vmware-horizon-view-best-practices-performance-study.pdf>



3.2.2 Citrix XenDesktop for VDI

Citrix XenApp is an application virtualization product for connecting to corporate applications from a wide range of computer systems and mobile devices. XenApp can host applications on central servers for remote user interaction or stream and deliver those applications to user devices for local execution.

Citrix XenDesktop is a suite of desktop virtualization products. Administrators can provision images using Citrix Provisioning Services (PVS), which provides single instance image management of XenApp and XenDesktop VMs, or Citrix Machine Creation Services (MCS), which uses linked clones to configure, start, stop, and delete VMs. Unlike VMware Horizon, XenDesktop can run on Citrix XenServer, VMware ESXi, or Microsoft Hyper-V.

Deploying Citrix XenApp on vSphere leverages the benefits of virtualization. Using TrueNAS for the data pool hosting the XenApp VM and for the storage for the PVS or MCS server leverages the thin-provisioning and data integrity features of TrueNAS. As of XenDesktop 7, it is possible to use MCS for XenApp load. TrueNAS supports both PVS and MCS.

iXsystems recommends these best practices when deploying XenDesktop as a VDI:

- Use TrueNAS 11.3 or later.
- Use TrueNAS for the VMs, the PVS, the MCS, and the Microsoft SQL Server used by the XenDesktop Controller. Do not use storage on the PVS server or place it inside the VMs.
- Provision storage to the PVS server using iSCSI.
- Provision storage to the MCS server using NFS.
- Use sync=always for both the MCS and the PVS.
- Provision storage to the Microsoft SQL Server from TrueNAS using SMB or iSCSI.
- Place the Citrix write cache on TrueNAS so it can benefit from the ARC, L2ARC, ZIL, and OpenZFS data integrity. Storing the write cache on TrueNAS also ensures its high availability.
- Review the Citrix Best Practices for XenDesktop knowledge base article⁴.

⁴ <https://support.citrix.com/article/CTX132799>

Note that all traffic between the vDisk and the client passes through the PVS, regardless of where the vDisk resides. By using TrueNAS, the OpenZFS file system will improve vDisk deployment efficiency. Since a vDisk is shared by many clients, TrueNAS will cache the file I/O operations so that subsequent clients will not require a disk read to perform similar operations. Read data will stay in the cache until it is flushed to make space for newer data.

Disk storage management is important because a PVS can have many vDisks stored on it and each disk can be several gigabytes in size. Use the TrueNAS GUI to view the amount of storage used by the PVS server.

VDI can be performance intensive in some situations. Watch performance and IOPS and contact an iXsystems specialist if these get excessive or clients experience degradation.

3.3 Server/Application VMs

Running multiple operating systems and multiple applications on the same physical server reduces costs while increasing the efficiency and utilization of existing hardware. TrueNAS enables VMs to benefit from centralized storage and the data integrity capabilities of OpenZFS.

iXsystems recommends the following best practices when using TrueNAS as the storage backend to server virtualization:

- Use TrueNAS 9.3 or later.
- When using multiple TrueNAS systems, create one data pool per TrueNAS array.
- Provision the datapool using iSCSI and set sync=always on the pool.
- If snapshots are used to protect VMs, use the TrueNAS VMware-Snapshot utility to schedule snapshots which can be kept locally or replicated to another system.
- Review the VMware Best Practices white paper⁵.

3.4 VMotion

VMware vMotion is the term used to describe a live migration where an entire running virtual machine is moved from one physical server to another without downtime. The virtual machine retains its network identity, connections, active memory, and precise execution state, ensuring a seamless migration. Since TrueNAS has a connection to every vSphere host, the switch from running on the source vSphere host to the destination vSphere host does not affect the user.

VMware Storage vMotion describes the live migration of virtual machine disk files within and across TrueNAS storage arrays with no downtime or disruption in service. When the source and destination TrueNAS storage arrays are connected to the vSphere host, there is no array-based acceleration for the data movement. Instead, the vSphere host reads the data from the source TrueNAS array and then writes it to the destination TrueNAS array. After all the data is moved, the original data is removed from the source TrueNAS array.

Long Distance vMotion describes the transfer of a VM from one datacenter to another through a two-step process. First, Storage vMotion transmits the disk version of the VM. Then, Long Distance vMotion transfers the current state of the machine to the new location. Long Distance vMotion can be used for any non-disruptive use case, including datacenter load balancing, maintenance, and disaster avoidance. Customers generally expect these use cases to include moving the actual storage in full transit.

When using TrueNAS on both sides, each must be active and accessible by both the source and destination VMware ESXi servers at all times. Latency between the source and the destination vSphere client should be no more than 5ms round trip time.

⁵ <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmware-perfbest-practices-vsphere6-0-white-paper.pdf>

3.5 VAAI

VMware vSphere Storage APIs Array Integration (VAAI) enables vSphere to offload specific VM and storage management operations, such as large data moves, to TrueNAS. Assisted by the storage hardware, vSphere performs these operations faster and consumes less CPU, memory, and storage fabric bandwidth. VAAI is supported by iSCSI.

Since TrueNAS is a thin-provisioned array, allocated data is created from a common storage pool. Consider a TrueNAS system with 20 TB of space hosting twenty virtual machines. Each machine has been given 1 TB of disk space for future use, but is currently only using 100 GB of that space. With ordinary thick provisioning, the entire allocated space is used, regardless of whether any actual data is stored there. The twenty VMs take up the entire 20 TB of available space. But each of the twenty VMs is really only using 100 GB, a total of only 1 TB. Thin provisioning allows TrueNAS to allocate only the space needed for actual data storage, leaving 19 TB free. In addition, TrueNAS coordinates VMware snapshots with TrueNAS snapshots, further reducing the storage consumed by VMware snapshots.

Data pools provisioned from TrueNAS storage will show up as thin in vCenter. TrueNAS will reclaim blocks of space when a VM is deleted or moved to a different array. When 80% of the available storage in a TrueNAS array is used, a warning will be displayed in vCenter. When space runs out, an error message is displayed in vCenter and affected VMs are paused so that the administrator can add storage or migrate the VMs to a different array.

VAAI integration offloads the following storage tasks from the ESXi server to the TrueNAS storage array:

- **Block Zero - Write Same Zero:** offloads repetitive write commands resulting in the ESXi server issuing fewer I/O requests for common tasks.
- **XCOPY (Full Copy):** offloads copying, snapshotting, and movement of blocks to TrueNAS, which enables faster VM deployments, cloning, snapshots and Storage vMotion.
- **Atomic test and set:** allows TrueNAS to control the locking of affected blocks instead of requiring ESXi to lock a full LUN. This enables more VMs per datastore by holding I/O for only the affected VMs.
- **UNMAP:** TrueNAS coordinates unallocated space with ESXi, limiting the need to add more storage.
- **Thin Provision Warn and Stun:** ESXi is aware of the thin-provisioned nature of TrueNAS and will issue a warning if more than 80% of the TrueNAS space is used and will issue an error message and pause affected VMs when 100% of the space is used. This prevents the VM from getting an I/O error when TrueNAS runs out of space.
- **NFS Space Reclaim:** Similar to UNMAP for iSCSI, it will unlink files clearing unused VMFS space in the NFS datastore.
- **LUN Reporting:** allows a hypervisor to query the NAS to determine whether a LUN is using thin provisioning.
- **Fast File Clone:** Allows the creation of linked-clones to be offloaded to TrueNAS.

3.6 Best Practices for VAAI

- Use TrueNAS version 11.3 or later for iSCSI.
- Provision the datastore using iSCSI (block storage) using device-based extents (zvols).
- Use iSCSI multipathing (MPIO).
- Depending upon the importance of the data, consider sync=always for iSCSI transactions.
- Connect the TrueNAS storage controller, or the two storage controllers of a TrueNAS HA configuration, to every server in a VMware Cluster.
- Set the VM timeout to 300 seconds on the guest operating system.
- For highly available storage, use a TrueNAS X10-HA, TrueNAS X20-HA, Z30-HA, Z35-HA, or Z50.
- For best performance, configure the disks into at least 5 mirrored vdevs.
- Use thick provisioned eager zeroed on vSphere 6 and higher if the guest operating system supports it. Otherwise, use a sparse volume to ensure thin provisioning.
- Verify throughput by running performance benchmarks on VMware and TrueNAS.
- Do not enable deduplication when using linked clones.
- Create a tunable with the Variable name of `vfs.zfs.l2arc_noprefetch`, a Value of 1, and a Type of Sysctl.
- Isolate the storage traffic from other networking traffic using VLANs, network segmentation, switches dedicated for iSCSI traffic, and jumbo frames with an MTU of 9000.
- Use iSCSI MPIO instead of a link aggregation protocol, such as LACP.
- Do not use VMware iSCSI port binding. Instead, each VMkernel adapter used for iSCSI traffic should be on a different IP subnet that matches the TrueNAS network configuration.
- Use ESXi Round Robin PSP (Path Selection Plug-in) for MPIO and adjust the Round Robin IOPS limit to 8.
- Configure the VM using the VMware paravirtual SCSI controller and use a VMXNET3 network adapter.
- For a Windows client or server, configure a centralized swap that resides on its own data pool.
- Install the VMware client tools.

4 TrueNAS® Overview and Portfolio

iXsystems designs and provides many TrueNAS Enterprise unified storage arrays for a variety of use cases. A TrueNAS array consolidates storage by supporting multiple file and block shares. TrueNAS Enterprise shares the same code base as TrueNAS CORE, so it benefits from the extensive feedback and testing of the TrueNAS community. iXsystems also designs the hardware of the array specifically for hybrid storage and TrueNAS.

TrueNAS is powered by OpenZFS and supports VAAI for iSCSI, making it the ideal platform for VMware object storage. TrueNAS leverages TrueCache™, which combines high-speed read and write caches for optimal performance.

Two TrueNAS families are available: hybrid and all-flash. The hybrid family uses flash for performance and hard disk drives for capacity. The all-flash family combines flash for performance and solid state disks (SSDs) for capacity. Each TrueNAS array has a modular design, is field replaceable, and has redundant power and cooling. All versions come with next business day controller replacement and a hybrid model comes in either a single- or dual- controller (HA) configuration. Single-controller configurations can be upgraded to a dual-controller configuration to ensure reliability.

The hybrid TrueNAS version supports flash for cache, SAS SSDs, and 7.2K, 10K, or 15K RPM hard disk drives for capacity. As storage requirements grow, additional drives and expansion shelves can be added. The TrueNAS “X” and “M” product families also have all-flash options in a dual-controller configuration for the utmost in reliability for high performance workloads.

5 Conclusion

TrueNAS is designed for virtualization and is an ideal platform for any organization that creates and uses multiple VMs. TrueNAS is a unified storage array that includes snapshots, storage optimization, multiple protocols, and replication that keeps VMs safe. There’s no question that TrueNAS is the best value in storage, period.

Don’t take our word for it. DCIG, a leading analyst site, ranked TrueNAS “Excellent” with “Best in Class” hardware⁶. If you want to talk about your virtualization needs, call one of our Account Representatives at 1.855.GREP4.IX, visit our web page at www.truenas.com, or email us at info@ixsystems.com. We look forward to hearing from you!

⁶ <https://www.ixsystems.com/blog/ixsystems-truenas-tops-latest-dcig-buyers-guide/>