



---

## **TRUENAS FOR VMWARE**

*Storage and Servers Driven by Open Source*

[marketing@ixsystems.com](mailto:marketing@ixsystems.com)

# CONTENTS

- 1 Executive Summary
- 2 History of iXsystems
  - 2.1 Open Source Background
  - 2.2 FreeNAS Overview
- 3 vSphere
  - 3.1 Introduction
  - 3.2 VMware HA
    - 3.2.1 Overview
  - 3.3 Virtual Desktops
    - 3.3.1 VMware Horizon View for VDI
    - 3.3.2 Virtualized Citrix XenApp for Virtual Desktops
    - 3.3.3 Citrix XenDesktop for VDI
  - 3.4 Server/Application VMs
    - 3.4.1 TrueNAS with Server Virtualization Best Practices
  - 3.5 vMotion
    - 3.5.1 vMotion
    - 3.5.2 Storage vMotion
    - 3.5.3 vMotion over distance
  - 3.6 VAAI
  - 3.7 Snapshot Coordination
  - 3.8 Best Practices For Your Datastore
- 4 TrueNAS
  - 4.1 Overview
  - 4.2 Design
  - 4.3 Portfolio
- 5 Conclusion
- 6 Acknowledgements

# 1 EXECUTIVE SUMMARY

Using a hypervisor means that you can 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, but if you aren't using TrueNAS you are doing nothing to protect your virtualized environment from the high percentage of silent data corruption that studies say exists.

TrueNAS is a storage array, based on the ZFS file system. The most important feature sets in ZFS are those designed to ensure the integrity of data. It uses features like end to end checksumming and a 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. Data integrity is a high priority in ZFS because research has shown that none of the currently widespread general purpose file systems commercial filesystems, nor hardware RAID provide sufficient protection against such problems. ZFS was also designed to overcome the limitations of previous file systems, RAID controllers, and volume managers, and uses advanced read and write caching algorithms that allow the application 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 now arguably the most precious commodity companies possess. TrueNAS offers unified and scalable storage solutions that can protect that data, simplify its management, reduce the operational costs, and optimize performance of a virtualized environment. It unifies SAN and NAS in one array and has integration with VMWare's VAAI, ZFS snapshots, thin provisioning, and VMware Ready certification. Easy to deploy and maintain, TrueNAS helps to ensure that critical business data remains available; through innovative double-fault protection and its use of OpenZFS, storage is protected from disk, network, and storage node faults.

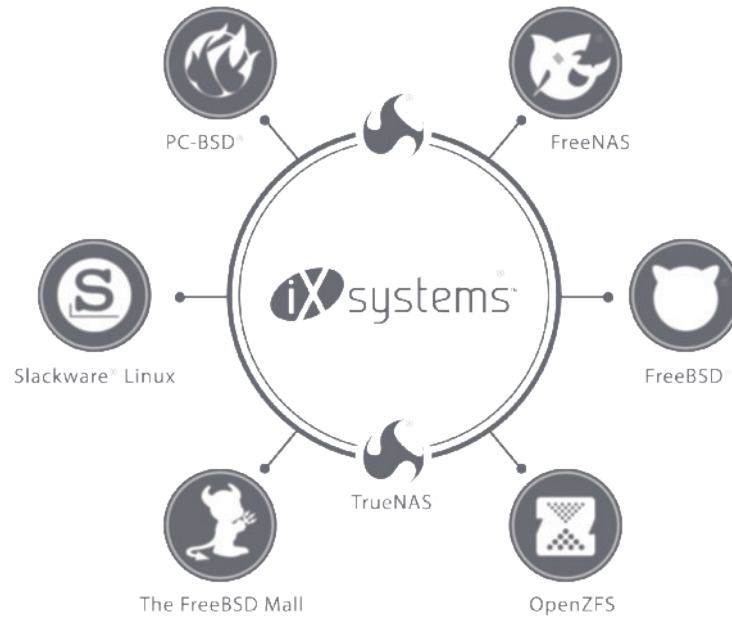
When coupled with VMware, TrueNAS will coordinate the snapshots such that VMware snapshots are converted to ZFS snapshots. This enables you to use them as stable local or remote resurrection points. Also, with its dynamic caching and storage optimization technology, TrueNAS eliminates the need to rely on multiple spindles of power hungry or short-stroked disks to get the IOPS for your VM deployment. The powerful features of variable block sizes, high performance storage optimization, snapshots, replication, and flash I/O acceleration adds to TrueNAS value. Deploying a TrueNAS enterprise storage array reduces overall cost and management overhead per VM.

## 2 HISTORY OF IXSYSTEMS & TRUENAS

Nestled in the heart of Silicon Valley, we have been committed to serving technology needs with a focus on Open Source since our beginnings as Telenet Systems in 1996. We became a part of BSDi in 2000 and have been iXsystems, Inc. since 2002. From our inception onward, we've been perfecting the craft of making quality, hand-tailored servers and storage solutions, and our dedication to a superior customer experience has been unwavering.

### 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 PC-BSD, FreeNAS, and TrueNAS 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.



Our CTO is Co-Founder of the FreeBSD Operating System, and we employ a long list of FreeBSD and Open Source project committers, all of whom came to iXsystems because of their passion for Open Source. It has allowed us to be successful, and in turn, allows us to help our clients do the same.

## 2.2 FREENAS OVERVIEW

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

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

## 3 VSPHERE

### 3.1 INTRODUCTION

vSphere is VMware's brand name for their suite of tools for cloud computing that use the ESX/ESXi hypervisor. The current release is vSphere 6.0, which included ESXi 6.0. TrueNAS supports vSphere 4.1 and later.

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

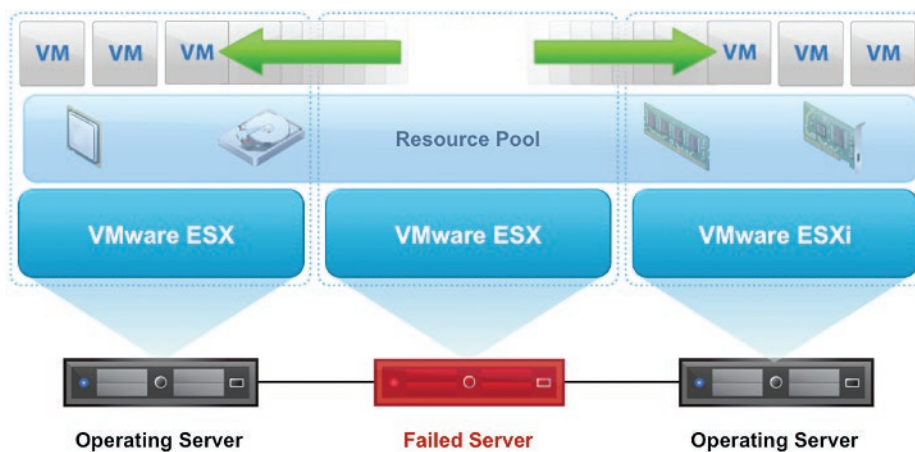
Log.in Consultants in their 2014 Project Virtual Reality Check surveyed nearly 1,000 enterprise customers and

found that customers are not only virtualizing their server applications, but almost 94% of them are also virtualizing Server-Based Computing (SBC) deployments (such as XenApp). Over half of these customers virtualize using vSphere, over one third use Citrix XenServer and a little more than ten percent use Hyper-V to deploy their virtual desktops.

## 3.2 VMWARE HA

### 3.2.1 OVERVIEW

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.



## 3.3 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 comes 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.

Many customers have virtualized their server applications and are now virtualizing Server-Based Computing (SBC) deployments of Citrix XenApp, Microsoft RDS, or are deploying virtual desktops hosted using VMware Horizon View or Citrix XenDesktop.

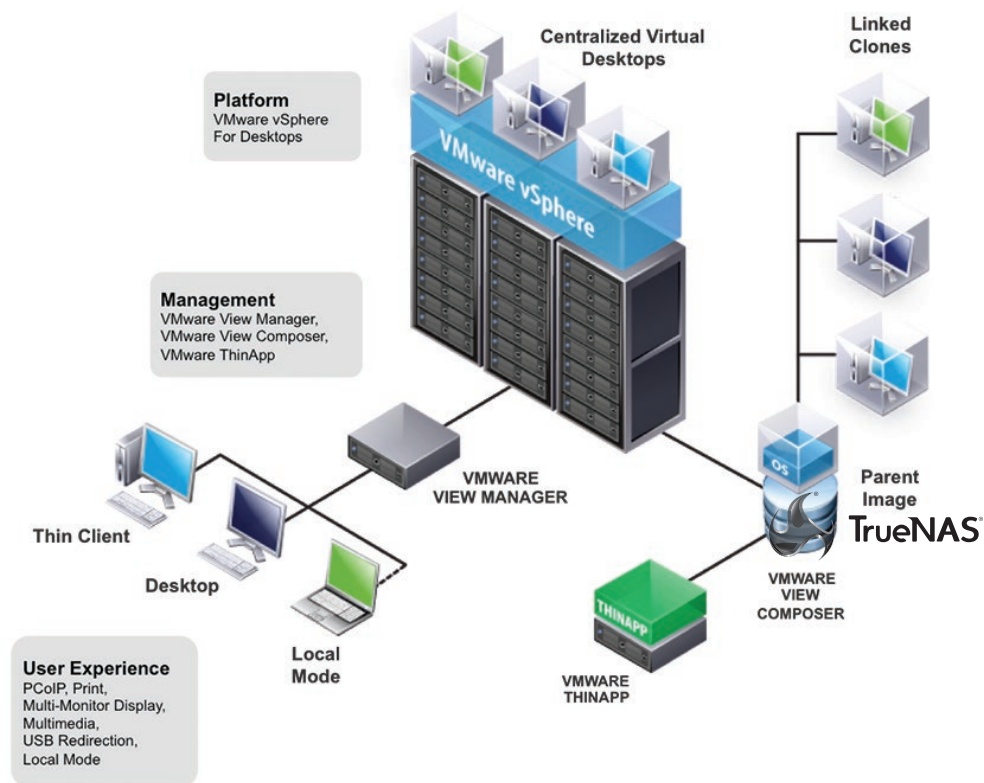
### 3.3.1 VMWARE HORIZON VIEW FOR VDI

VMware Horizon View is also known as VMware View. It 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. The VMware View product has a number of components that may be used, but are beyond the scope of this paper:

- VMware vCenter Server (management of virtualization environment)

- View Composer (advanced View management, with automation and cloning)
- View Manager (administration of the View Environment)
- View Client (communication between View and the desktop OS)
- VMware ThinApp (application virtualization)
- View Persona Management (user profile management)
- vShield Endpoint (offloaded desktop antivirus)

### 3.3.1.1 TRUENAS WITH HORIZON VIEW BEST PRACTICES



- Use TrueNAS 9.3 or later, which has VAAI. The use of VAAI will increase the number of VDI VMs supported by your server, since some of the work necessary does not occupy server resources but has been off-loaded to TrueNAS.
- VMware View relies on a SQL Server that is running on one of the servers in the vSphere cluster. It would typically be configured to use server storage for its data files (logs and databases), iXsystems recommends that TrueNAS be used instead. That way View will leverage the data integrity capabilities of ZFS.
- If you are running a BSD-based client, use VMware View Open Client -- it uses rdesktop, which is in FreeBSD
- Enable desktop and email searching. Many Horizon View best practices say you should disable this, but the read caching in TrueNAS ensures that searching is nearly instantaneous.
- Provision the datapools and SQL Server storage using iSCSI. When you define each iSCSI dataset use `sync=always` to avoid problems of data loss if your TrueNAS is lost
- If you use snapshots to protect your VDI VMs, then use the TrueNAS snapshot manager to create manual or automatic VMware snapshots
- Review the VMware Horizon View Best practices [here](#).

### 3.3.2 VIRTUALIZED CITRIX XENAPP FOR VIRTUAL DESKTOPS UNDER VMWARE

Citrix XenApp is an application virtualization product that allows users to connect to their corporate applications from a wide range of computer systems and mobile devices. XenApp can host applications on central servers and allow users to interact with them remotely or stream and deliver them to user devices for local execution.

Customers are virtualizing deployments of Citrix XenApp on vSphere to leverage the benefits of virtualization. By using TrueNAS for the datapool hosting the XenApp VM and for the storage for the PVS or MCS server they leverage the thin-provisioning and data integrity of TrueNAS.

Before XenDesktop 7, it wasn't possible to do MCS for XenApp load. Since XenDesktop 7 it's actually possible to use Machine Creation Services for XenApp. TrueNAS supports both PVS and MCS see the XenDesktop best practices below.

### 3.3.3 CITRIX XENDESKTOP FOR VDI

Citrix XenDesktop is a suite of desktop virtualization products from software provider Citrix Systems. Administrators can provision images using either Provisioning Services (PVS), which provides centralized disk management, or Machine Creation Services (MCS), which uses linked clones. Then, admins can deliver desktops and apps to users through five different FlexCast methods.

Unlike VMware View, XenDesktop can run on Citrix XenServer, VMware ESXi, or Microsoft Hyper-V. This paper only covers VMware.

#### 3.3.3.1 TRUENAS WITH XENDESKTOP BEST PRACTICES

- Use TrueNAS 9.3 or later
- Use TrueNAS for the VMs, the PVS and MCS server.
  - Provision storage to the MCS server using NFS
  - Provision storage to the PVS server using iSCSI and use `sync=always`
    - Do not use storage on the PVS server or place it inside the VMs, instead use TrueNAS.
      - All traffic between the vDisk and the client passes through the PVS machine regardless of where the vDisk resides. By using TrueNAS, the ZFS file system will improve vDisk deployment efficiency. Since a vDisk is shared by many clients, TrueNAS will cache the file IO operations so that subsequent clients will not require a disk read to perform similar operations. The read data will stay in the cache until it is flushed to make space for newer data.
    - Disk storage management is very important because a PVS server 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.
    - The XenDesktop Controller uses Microsoft SQL Server - provision storage to it from TrueNAS and use CIFS/SMB or iSCSI
- Place the Citrix write cache on TrueNAS. The write cache will benefit from the ARC, L2ARC, the ZIL, and the ZFS data integrity. Storing the write cache on TrueNAS ensures its high availability.
- VDI can be performance intensive in some situations. Watch performance and IOPS and contact an iXsystems specialist if those get excessive or client user experience degrades.
- Review the Citrix best practices for XenDesktop [here](#).

## 3.4 SERVER/APPLICATION VMS

Running multiple operating systems and multiple applications on the same physical server lets you reduce costs while increasing the efficiency and utilization of your existing x86 hardware. Using TrueNAS enables these VMs to benefit from centralized storage and the data integrity capabilities of ZFS.


### 3.4.1 TRUENAS WITH SERVER VIRTUALIZATION BEST PRACTICES

- Use TrueNAS 9.3 or later
- Allocate one datapool to TrueNAS
  - If you have more than one TrueNAS then don't span a datapool across every TrueNAS array but create one datapool for every TrueNAS array
- Provision your datapool using iSCSI
  - iSCSI does not use synchronous writes, so your VMs may be at risk. Enable synchronous writes. Set `sync=always` on the dataset.
- If you use snapshots to protect your VM, then use the TrueNAS snapshot manager to create manual or automatic VMware snapshots, which can be kept locally or replicated
- VMware has multiple best practices on its web site. Review them.

## 3.5 VMOTION

### 3.5.1 VMOTION

VMware vSphere live migration (aka vMotion) allows you to move an entire running virtual machine from one physical server to another without downtime. The virtual machine retains its network identity and connections, ensuring a seamless migration process. TrueNAS has a connection to every vSphere host. The transfer is of the virtual machine's active memory and precise execution state, allowing the virtual machine to switch from running on the source vSphere host to the destination vSphere host, without affecting the user.



### 3.5.2 STORAGE VMOTION

VMware Storage vMotion provides live migration of virtual machine disk files within and across TrueNAS storage arrays with no downtime or disruption in service. The source and destination TrueNAS storage arrays are connected to the vSphere host. There is no off-load of the data movement, 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.

### 3.5.3 VMOTION OVER DISTANCE

Long-distance vMotion enables the transfer of a virtual machine (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 many, many things – datacenter load balancing, maintenance, and disaster avoidance (among other things).

The key is that the use case demands that it is non-disruptive in all use cases – and customers generally expect that to include moving the actual storage in full transit.

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

## 3.6 VAAI

VMware's vStorage APIs for Array Integration (VAAI) enables vSphere to offload specific virtual machine and storage management operations to TrueNAS. With the storage hardware assistance, vSphere and your VMs perform these operations faster and consume less CPU, memory, and storage fabric bandwidth. VAAI is only supported with iSCSI.

TrueNAS is a thinly provisioned array, so the allocated data is created from a common pool of storage. For example, if you allocate 1 TB to 20 VMs and each VM uses 100GB, a thinly provisioned array will consume 1TB of the 20TB of physical storage is consumed, but a thickly provisioned storage array would consume all 20TB.

Datapools that include storage provisioned from TrueNAS will show up as thin in vCenter. Additionally, TrueNAS will reclaim blocks of space when a VM is moved to a different array or is deleted. 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, giving admins time to add storage or migrate the VMs to a different array.

TrueNAS also coordinates VMware snapshots with TrueNAS snapshots, which reduces the storage consumed by VMware snapshots. VAAI integration offloads the ESXi server and has the TrueNAS array perform the function:

- Block Zero - Write Same Zero
  - Offloads common and repetitive write commands resulting in the ESXi server issuing fewer I/Os for common tasks
- XCOPY (Full Copy)
  - Offloads the copying, snapshotting and movement of blocks to TrueNAS, which enables faster VM deployments, cloning, snapshots and storage vMotion.
- Atomic test and set
  - Lets the TrueNAS control the locking of affected blocks instead of ESXi needing to lock a full LUN, enabling more VMs per datastore by holding IO for only the affected VMs
- UNMAP
  - TrueNAS coordinates with ESXi to understand what allocated space is no longer needed by ESXi, limiting the need to add more storage to TrueNAS

- Thin Provision Warn and Stun
  - ESXi is aware of the thinly 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 IO error when TrueNAS runs out of space.

## 3.7 SNAPSHOT COORDINATION

TrueNAS allows you to coordinate ZFS snapshots with VMware. You use the TrueNAS GUI to create a manual or automatic VMware snapshot. Once this type of snapshot is created, TrueNAS will tell vSphere to snapshot any running VMware virtual machines before taking a ZFS snapshot of the dataset or zvol backing that VMware datastore. TrueNAS then tells vSphere to delete them. The VM snapshots are retained as ZFS snapshot and can be used as stable resurrection points in that snapshot. They are also listed in the "Snapshots" tab of the TrueNAS GUI with a name derived from the VM name.

By leveraging ZFS snapshots, TrueNAS ensures that no space is initially required in the datapool to contain the snapshot. A VMware snapshot will occupy space in the datapool. When TrueNAS tells vSphere to delete the VMware snapshots, VAAI will be used to tell TrueNAS that the space can be reclaimed.

## 3.8 BEST PRACTICES FOR YOUR DATASTORE

- Use TrueNAS version 9.3 or later
- Provision your Datastore using iSCSI (block storage)
- You should always have a flash-based ZIL (SLOG)
- Use one network path for every iSCSI share
- Use iSCSI multipathing
- Use iSCSI shares for your datapool with "sync=always"
- Enable LZ4 compression for each iSCSI share used for the datapool
- Connect the TrueNAS storage controller (two storage controllers if you use a TrueNAS HA configuration) to every server in a VMware Cluster
- Set the VM timeout to 300 seconds
- To ensure highly available storage, use a Z20-HA, Z30-HA, Z35-HA or Z50
- Use at least 5 vdevs if you're concerned about performance
  - Use mirrored vdevs
- Use Eager-Zeroed Thick provisioning for the VMs
- Verify your implementation by running performance benchmarks on VMware and TrueNAS to validate you're getting the throughput expected
- If linked-clones are used then ensure deduplication is off
- Disable prefetching
  - `vfs.zfs.l2arc_noprefetch=1`
- Use a sparse volume (thin provisioning)
- Isolate the storage traffic from other networking traffic
  - Use VLANs
  - Use network segmentation
  - Use switches dedicated for iSCSI traffic
- Configure at least 2x10 GbE NICs (each NIC card in TrueNAS is dual ported) and bonded into a single channel using LACP
- Use jumbo frames (9000 bytes)
- Configure VMware iSCSI to work in port binding mode
- Configure your VM using the VMware paravirtual SCSI controller

- Use a VMXNET3 network adapter
- For Windows (client or server) configure a centralized swap that resides on its own datapool
- Install the VMware client tools

## 4 TRUENAS

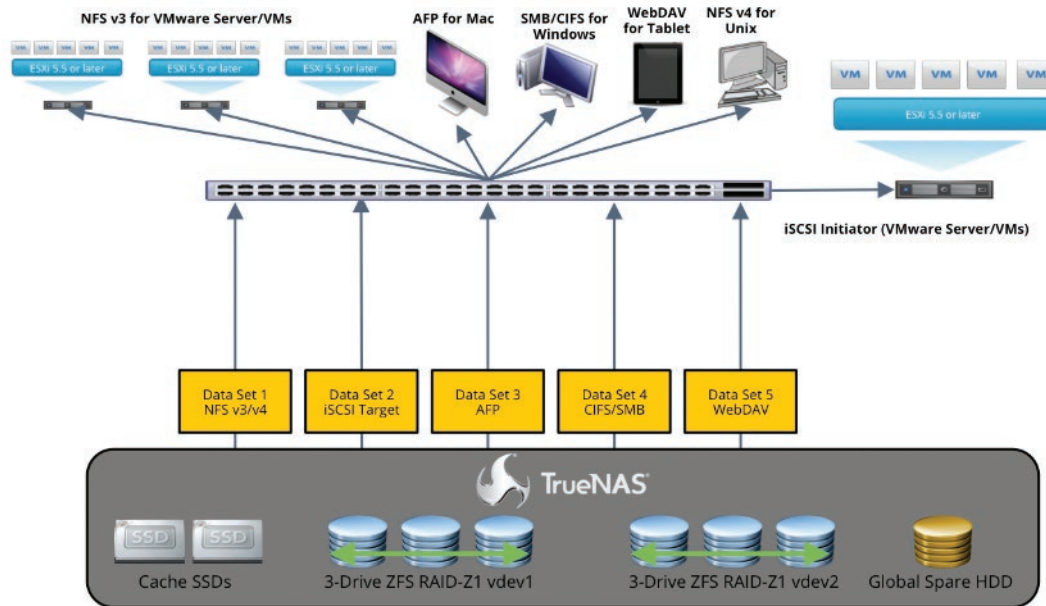
A TrueNAS array helps you consolidate storage that is dedicated for a single server or single application from your environment by supporting multiple file and block shares.

### 4.1 OVERVIEW

TrueNAS is a Unified Storage Array developed and designed by iXsystems. The TrueNAS software shares the same code base with FreeNAS, so it benefits from the extensive feedback and testing of the FreeNAS community. iXsystems also designs the hardware of the array specifically for Hybrid Storage and TrueNAS.

### 4.2 DESIGN

TrueNAS is an enterprise storage array that is powered by OpenZFS and exports 5 types of shares (NFS v3/v4, AFP, SMB/CIFS, WebDAV and iSCSI). iXsystems recommends iSCSI for VMware, however all protocols leverage TrueCache, which combines high-speed read and write caches in front of storage. When you use TrueNAS you get the high performance read/write of flash and the capacity of drives.



TrueNAS supports iSCSI and NFS making it ideal for storage of your VMware objects. VMware's VAAI is supported with iSCSI, which lets the array off-load work from vSphere. Finally, TrueNAS integrates with vSphere to take a VMware snapshot, which is then converted into a space efficient, replicable ZFS snapshot.

## 4.3 PORTFOLIO

All TrueNAS models are powered by OpenZFS. There are two models, a hybrid version and an all flash version. All versions use TrueCache. The hybrid version combines flash for performance and hard disk drives for capacity. The all flash version 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 the hybrid model comes in either a single- or dual-controller (HA) configuration. Single-controller configurations can be upgraded later to a dual-controller configuration to ensure reliability.

The hybrid TrueNAS versions support flash for cache and SAS SSDs, 7.2K, 10K, or 15K RPM hard disk drives for capacity. Additionally, you can grow as your storage requirements grow by adding drives/expansion shelves. The all-flash TrueNAS Z50 TrueFlash version comes exclusively in a dual-controller configuration for the utmost in reliability required by the highest performance workloads. It uses 15 SSDs, giving you the ultimate in performance and low latency that comes from its use of non-rotating media.

## 5 CONCLUSION

You can use all enterprise features of TrueNAS without having to pay extra. It is a unified storage array that includes snapshots, storage optimization, multiple protocols, and replication for less than most competitor's basic model price. There's no question that TrueNAS is hands-down the best value in storage, period. Don't take our word for it, see why DCIG, a leading analyst site, ranked TrueNAS "Excellent" with "Best in Class" hardware. Learn more when you by calling one of our consultative representatives at 1.855.GREP.4.IX, visiting our web page at [www.ixsystems.com/TrueNAS](http://www.ixsystems.com/TrueNAS), or emailing us at [sales@ixsystems.com](mailto:sales@ixsystems.com). We look forward to hearing from you.

## 6 ACKNOWLEDGEMENTS

The author would like to thank **Josh Paetzel** for his assistance.

*This White Paper is for informational purposes only. **IXSYSTEMS MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS WHITE PAPER.** iXsystems cannot be responsible for errors in typography or photography.*

*iXsystems, the iXsystems logo, TrueNAS, FreeNAS and the FreeNAS logo are registered trademarks of iXsystems. FreeBSD and the FreeBSD logo are registered trademarks of the FreeBSD Foundation. Microsoft and Windows are registered trademarks of Microsoft Corporation. UNIX is a registered trademark of The Open Group. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. iXsystems disclaims proprietary interest in the marks and names of others.*

*Information in this document is subject to change without notice.*