



WHITE PAPER

# TrueNAS® - Sharing Billions of Files

---



# CONTENTS

---

1 Executive Summary

2 History of iXsystems

3 Sharing Files

3.1 Direct Attached Storage File Sharing

3.2 Direct Attached Storage Issues

3.2.1 Complexities and Resource Inefficiencies

3.2.2 Cost Increases with lower Productivity

3.2.3 Poor Scalability

3.2.4 Limited Support of Multiple Applications

3.3 Storage Area Network File Sharing

3.4 Storage Area Network Issues

3.5 Network Attached Storage File Sharing

3.6 Consumer-Grade NAS Issues

3.7 TrueNAS - Sharing Billions of Files

3.7.1 The Value of using TrueNAS to Share Files

3.7.1.1 Share Files to Everyone

3.7.1.2 Keep Backups

3.7.1.3 Host and Stream Multimedia Files

3.7.1.4 Virtualize Applications

4 Conclusion



# 1 Executive Summary

---

Companies use data to improve their bottom line. Data is growing exponentially, and many companies are dealing with “big data” and need to handle data generated by multiple devices, such as smartphones, smart tablets, laptops, wearable technology, and other data sources from the Internet of Things (IoT). A lot of this data is shared locally and remotely and it would be difficult to find an organization that is not using file sharing in some shape or form. Not all file sharing services are created equal and utilizing the correct type of file sharing is crucial to success.

Organizations often begin sharing files using large numbers of direct attached storage (DAS) file servers. The problem with this strategy is that each DAS must be separately acquired, provisioned, and maintained. To reduce the associated storage costs, IT organizations often move to using fewer file servers built from storage area networks (SAN) or network attached storage (NAS). But as data grows, multiple solutions are deployed, each requiring management and maintenance.

Over time, IT is often forced to over-provision, and over-purchase, capacity or performance or use an all-flash array to meet an application’s needs. This drives up the cost of data storage solutions. This paper demonstrates how TrueNAS lets you grow capacity or performance independently.

This white paper is intended for CIOs, systems integrators, and systems and storage administrators. It provides an overview of file sharing with DAS, SAN, and NAS, an introduction to file sharing concepts, and how TrueNAS can be used as a NAS to provide unrivaled data integrity protection with its OpenZFS enterprise file system.

## 2 History of iXsystems

---

Located in the heart of Silicon Valley, iXsystems has been committed to serving technology needs with a focus on open source and enterprise hardware since our beginning in 2002. 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.ixsystems.com/reviews/>.

## 3 Sharing Files

---

### 3.1 Direct Attached Storage File Sharing

Direct attached storage (DAS) is storage connected to one computer and is not accessible to other computers. Deploying a new application requires that a new computer be acquired for the application and provisioned with enough storage to ensure that the application does not run out of space. Additionally, horsepower is shared between the storage and the application. This can degrade the application as well as the user experience.

## 3.2 Direct Attached Storage Issues

### 3.2.1 Complexities and Resource Inefficiencies

#### Complexity:

- **Access:** When a system needs to access a file that resides on another computer, the file must first be copied to that system. If changes are made to the file on the second system, it must be copied back.
- **Availability:** A DAS server does not provide High Availability (HA). Every file must be copied to a secondary DAS server that has the same configuration as the original and recopied when the file is changed on the original DAS server. In the event of an outage on the original DAS server, the second DAS server must be used. When the outage is over, all intervening file changes must be copied back to the original DAS. The time to copy these files, as well as the acquisition and maintenance of the second DAS server increases the cost of the file.
- **Growth of "islands of storage":** Every new application requires IT to either deploy a new DAS server, or add additional hardware, networking, or storage to existing servers to support additional applications or data types.

#### Resource Inefficiency:

- **Use of resources:** Using a server as a DAS requires the hardware and networking horsepower to handle files and applications, increasing the cost of the server.
- **Inefficient distribution of resources:** If you overbuy storage capacity for one DAS appliance and underbuy for another, there is no way to shift capacity from one system to another without costly downtime and hours of labor. With a NAS, the available storage space seamlessly goes where it is needed.
- **Scaling:** Since storage in the server cannot be scaled non-disruptively or the file system used by the server may restrict scalability, IT will over-provision storage to accommodate future growth. This results in poor storage resource utilization, which also increases the cost of the server.

### 3.2.2 Cost Increases with lower Productivity

#### Storage Management Costs:

- **Deployment costs:** Every DAS server must have an OS installed and updated to the latest patch/upgrade level. It must also be provisioned to hold the data.
- **Other costs:** Every DAS server has on-going needs for backup, archive, Disaster Recovery (DR), support, and security.
- **Proliferation of management tools:** IT can end up using different management GUIs or APIs for each DAS, causing proliferation of management tools.

### 3.2.3 Poor Scalability

A DAS cannot scale like a NAS and it is difficult to implement SLAs for multiple DAS systems. Scaling or protecting a DAS impacts business operations as the application and its data are not available during an outage.

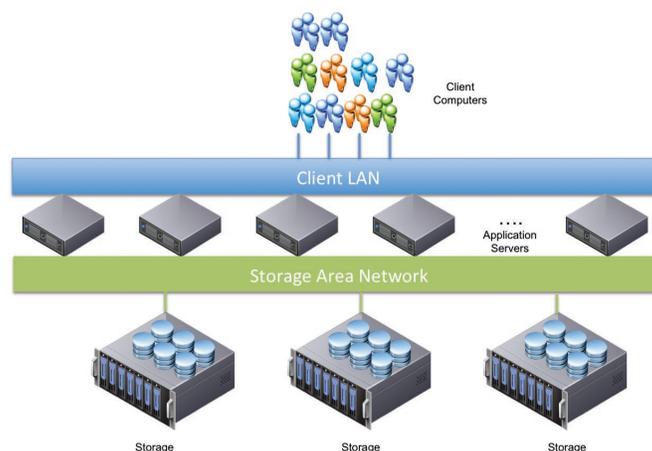
Scaling requires staff to power down the DAS and insert drives into the empty bays which impacts file availability. If there are no empty bays, a compatible JBOD and storage controller must be acquired.

Protecting data availability requires organizations to deploy backup plans built around data replication. Many organizations delay implementing these solutions because they are considered complex and expensive to set up and manage, especially if new storage or networking hardware is required. This hesitation increases the risk of data loss.

### 3.2.4 Limited Support of Multiple Applications

In mixed Windows and UNIX environments, each platform uses a different protocol and file system. A DAS is limited by its inability to support more than one protocol, OS, or file system. When an application is added or grows, the appropriate DAS server must be expanded or a new DAS server added. This leads to storage sprawl and hinders an organization's ability to accommodate application and data growth.

## 3.3 Storage Area Network File Sharing



A SAN is a high-performance network whose primary purpose is to enable storage devices to communicate with computer systems using block protocols. The communication and data transfer mechanism for a given deployment is commonly known as a *storage fabric*.

## 3.4 Storage Area Network Issues

A SAN uses a separate network and is generally not accessible by other devices through the local area network. A SAN does not provide file abstraction, only block-level operations. Common SAN protocols are iSCSI, Fibre Channel, and Fibre Channel over Ethernet (FCoE). A proprietary management protocol can deliver additional capabilities such as disk zoning, disk mapping, LUN masking, and fault management.

Since a SAN does not provide file abstraction, a server operating system creates a file system from the storage and then shares the file system to clients attached to the server. The block organization of a SAN makes it costly and complex to share data between business locations. To avoid this, many IT organizations create a duplicate copy of the data at the remote location, reducing deployment complexity but increasing storage costs.

Despite their differences, SAN and NAS are not mutually exclusive and can be combined in multi-protocol or unified storage arrays. A TrueNAS unified storage array offers both NAS file-level protocols and SAN block-level protocols.

## 3.5 Network Attached Storage File Sharing

A NAS system addresses the shortcomings of DAS. A NAS is a file storage device connected to a computer network that provides data access to a heterogeneous set of clients. Unlike an application server with direct attached storage, a NAS is specialized for serving files.

## 3.6 Consumer-Grade NAS Issues

Consumer NAS devices can have limitations in performance, reliability, scalability, and network throughput due to limitations in CPU horsepower, low memory configurations, or use of consumer-grade hardware. A consumer-grade NAS can present several issues and limitations, including:

### Performance:

- Many consumer-grade NAS devices have I/O bottlenecks at the network ports, the controllers, or the disk drives
- Using consumer-grade NAS devices for virtualization storage exacerbates the problem as the status of the storage might not be shared with the hypervisor or the hardware may not be sufficient to scale to I/O requests from many VMs.
- Bottlenecks are often caused by the inability of a consumer-grade NAS to have the CPU horsepower to handle all the NFS metadata commands.
- Many consumer-grade NAS devices see poor storage performance as a result of high I/O latency.
- Lack of optimization for the multiple I/O operations generated by backups and media production/distribution.

### Economics:

- Lower-end NAS devices have a much higher total cost of ownership (TCO) when considering downtime, upgradability, and lackluster performance.
- Many NAS devices do not have enterprise technology or features such as HA, ECC memory, non-disruptive upgrades, or enterprise-class drives. Many NAS systems use low-cost or consumer-grade networking hardware, so an intermittent failure can affect other systems that share the same physical network, requiring time to diagnose and fix the problem.
- Configuration of read, write, and access modes for file sharing can be complex when dealing with multiple operating systems. Some customers buy more NAS hardware to dedicate to a single operating environment in an attempt to compensate for this problem.

### Inability to increase availability:

- For many NAS systems, an increase in performance or moving to a high-availability environment is not supported or requires an outage and a forklift upgrade. During the upgrade, no one can access the file content.

### Expandability:

- Traditional and consumer NAS systems have a finite capacity that is often smaller than many environments need. If a customer fills up the NAS, their choice is to buy another NAS, which increases storage sprawl, or perform a forklift upgrade to a different NAS vendor, increasing costs.

All of these issues and limitations increase the storage cost of data. The next section describes how TrueNAS solves these issues.

## 3.7 TrueNAS - Sharing Billions of Files

TrueNAS uses multiple techniques to reduce the volume of data actually stored. This storage optimization happens while the data is in transit and reduces the size of I/Os that reach storage media, which in turn increases performance while conserving space for additional data. This keeps data growth under control and further improves the TCO of TrueNAS hybrid and all-flash arrays.



TrueNAS is award-winning, enterprise-grade unified storage. It is both a NAS and a SAN supporting a wide variety of file and block protocols. Built from enterprise components, it provides a modular and tool-less architecture. It uses the enterprise-grade OpenZFS file system for delivering self-healing bit rot mitigation, flexible snapshotting, and replication.

Most companies purchase data storage based on capacity and performance dictated by the needs of existing applications. As a result, businesses often end up with multiple classes of application-specific storage or “storage silos” including SAN, NAS, all-flash arrays, and DAS. The issues that arise from these storage silos are summarized in Hybrid or All-Flash: The Choice is Yours.

A TrueNAS storage array from iXsystems uses ZFS caching, a combination of RAM and flash-based cache. It delivers blazingly fast performance for file sharing, dynamically storing the most often-used files onto the fastest storage media. As a unified solution, TrueNAS supports multiple protocols that make it possible to consolidate storage silos. It scales capacity to 9 PB without performance impact.

Every TrueNAS system includes data replication to safely replicate vital data to remote locations, such as a DR data center. Other vendors charge extra for this feature, increasing the TCO for data storage. By balancing price and performance, TrueNAS integrates as a core component in physical or virtual server DR plans. With TrueNAS in place, you can sleep well at night.

### 3.7.1 The Value of using TrueNAS to Share Files

TrueNAS allows multiple computers to share file data via a network connection using multiple file protocols. With TrueNAS, the data is presented to other servers across the network, while the storage management task is performed by the TrueNAS storage array.





TrueNAS protects files by calculating a checksum as every disk block is written. When a block is read, the checksum is recalculated and any corrupt data blocks are repaired from a redundant copy whenever the checksum does not match. A traditional file system does not use checksums and will instead return any corrupted data when it is read.

Using TrueNAS for file sharing provides these benefits:

#### **Improved Reliability:**

- The Self-Healing file system guarantees data integrity for every file. All reads return only verified data and any corrupted data is automatically repaired from a replica.

#### **Enhanced Supportability:**

- Award-winning US-based 24/7 worldwide support.
- A tool-less and modular design to lower service costs.

#### **Increased Availability:**

- A unified file system that supports Apple, Windows, and Unix systems.
- Upgradable to a redundant storage architecture to ensure that clients can still access data in the event of an outage.
- Redundancy and HA capabilities that are not present in most DAS solutions.
- Administrator-less recovery on various Windows-based systems that integrates TrueNAS with the volume shadow copy feature used by Windows Explorer "Previous Versions".
- WebDAV for increased mobile device support.

#### **Performance Boost:**

- Flash-assisted hybrid architecture for maximizing the performance of every I/O operation. A combination of RAM, flash memory, and hard disks provide the performance of flash with the capacity and economics of hard disks.

#### **Lowered Risk:**

- Certified by Citrix, VMware, and Veeam.
- Automated recovery points for VMs and other business-critical applications makes TrueNAS a core component of a DR plan.
- Non-disruptive software and hardware upgrades are more reliable and robust than DAS.
- Integration with enterprise directory services.

#### **Reduced Storage Cost:**

- Utilization of IP networks for file-level access reduces the price of entry for access to shared storage.
- Utilization of Fibre Channel infrastructure reduces costs when replacing legacy storage.
- TrueNAS delivers enterprise features like replication, thin-provisioning, snapshots, clones, in-line compression, and deduplication at no additional costs or add-on licenses.

The OpenZFS file system used in TrueNAS features many of the same, if not more than, enterprise-grade data integrity safeguards as most DAS and SAN systems, regardless of price. Here are four use cases that showcase the value of TrueNAS.

### **3.7.1.1 Share Files to Everyone**

Sharing is what TrueNAS does best. Every major operating system, application, and hypervisor is supported. With its highly granular quota and reservation tools, it is easy to scale TrueNAS as more space is required. Allocating storage ahead of availability is no risk and the use of a shared storage pool makes thick provisioning simple.

TrueNAS brings an enterprise storage solution that can be used in a remote office or a data center to lower storage expenses. TrueNAS includes enterprise capabilities such as snapshots, replication, encryption, and high availability, at a price lower than many competitors.

### **3.7.1.2 Keep Backups**

Data is critical to any business. The requirement to store and protect data continues to grow rapidly. TrueNAS includes automatic and instant point-in-time snapshots which can be used to build a backup or DR environment. Protect months' worth of data locally and remotely and reduce risk while lowering server, network, and storage costs.

A TrueNAS backup repository can be managed from across the room or across the globe, eliminating the need for remote office IT staff. Increase capacity as your business grows and reduce storage infrastructure cost and risk.

### **3.7.1.3 Host and Stream Multimedia Files**

Media and entertainment organizations face multiple storage challenges:

- Growth of Ultra High Definition (UHD) content, higher frame rates, more audio channels, and multiple transcoded versions of digital assets increases storage consumption. Storage products must be flexible and expandable.
- I/O performance and latency concerns from real-time videos and the need to produce content at a moment's notice increase pressure on storage performance.
- The use of global teams requires the storage to support a variety of operating systems and remote management.

Using TrueNAS to host and stream multimedia files addresses all these challenges, helping your business stay competitive. It brings an enterprise storage solution to media and entertainment applications. Replacing DAS with a TrueNAS enterprise storage array delivers safe, secure, unified, and scalable storage that can be used with multiple multimedia applications.

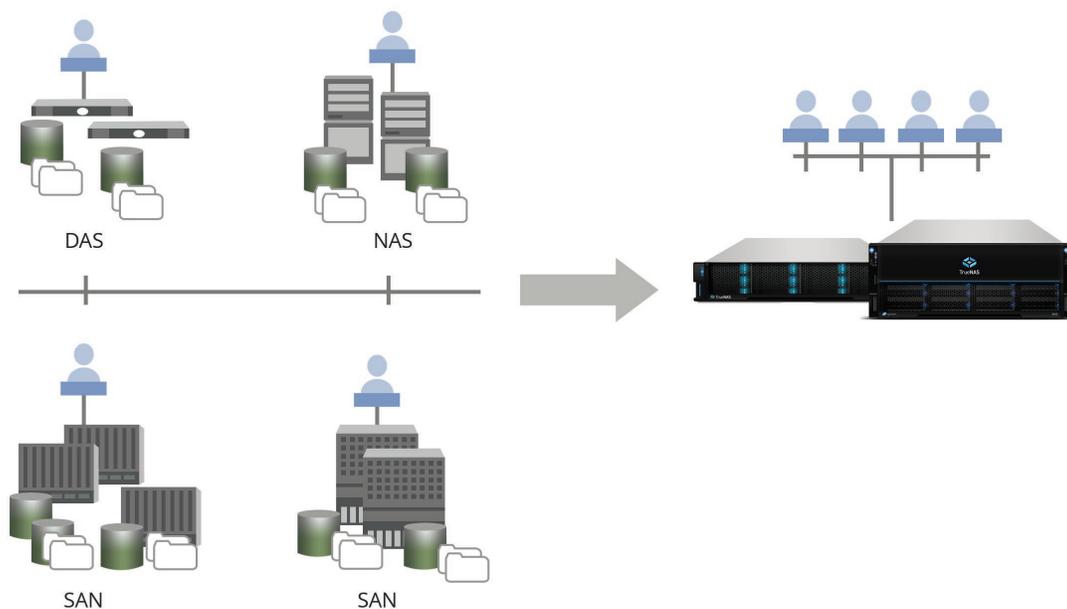
### **3.7.1.4 Virtualize Applications**

TrueNAS supports enterprise virtualization needs. It scales capacity to 9 PBs, unifies block, file, and object protocols, and scales bandwidth to over 40 Gb/sec. It integrates with all major virtual machine environments, enabling files of any size to be shared to hundreds or thousands of users. Run more operating system environments on a single host from a single, hassle-free array.

TrueNAS provides integrated support with VMware VAAI/Block, VAAI/NAS, and VMware snapshots, Veeam, Microsoft CSV, ODX, and VSS, as well as a vCenter plugin. TrueNAS can create instant crash-consistent snapshots of any VMware VM, allowing you to replicate a VM and restart it at the remote location.

TrueNAS has been certified by Citrix and VMware. All hypervisors from Citrix, VMware, Microsoft, and others are supported. TrueNAS provides integrated support with VMware VAAI/Block, VAAI/NAS, VMware snapshots, Veeam, Microsoft CSV, ODX, and VSS, as well as a vCenter plugin. TrueNAS can create instant crash-consistent snapshots of any VMware VM, allowing you to replicate a VM and restart it at the remote location.

TrueNAS reduces the complexity and inefficiency associated with DAS-based hardware. TrueNAS provides highly available, high performance, and feature-rich storage for a wide variety of enterprise and mission critical business applications, all built on top of a best-in-breed file system that guarantees data stays pristine and secure.



## 4 Conclusion

A business needs a fast and reliable solution to protect and share files. When TrueNAS is used for file sharing, its self-healing file system automatically repairs shared files after a corruption, such as bit rot. TrueNAS provides unified, enterprise-grade, fast performance for file read and write operations across all block or file protocols.

With the ability to share petabytes of data to thousands of VMs, there's no question that TrueNAS is hands down the best value in storage. Don't take our word for it, [see why DCIG](#), a leading analyst site, ranked TrueNAS "Excellent" with "Best in Class" hardware.