

Technical Review

iXsystems TrueNAS

Date: January 2020 Author: Tony Palmer, Senior Validation Analyst

Abstract

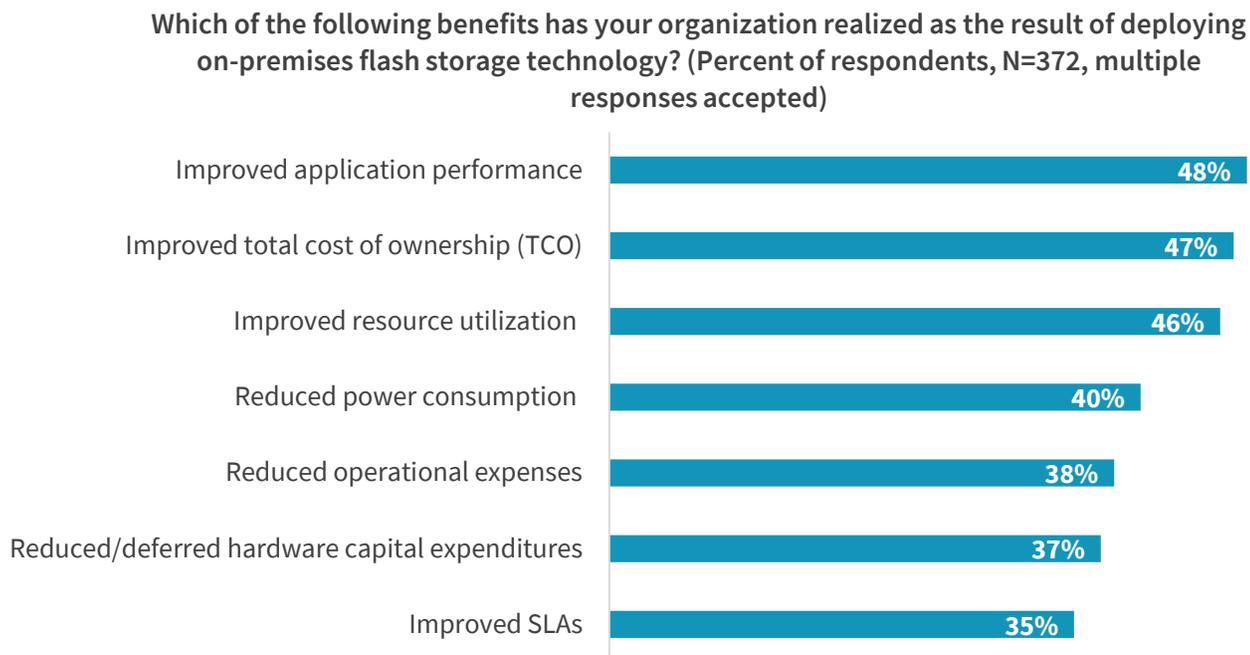
This ESG Technical Review documents iXsystems TrueNAS' storage performance, flexibility, and integration for the evolving on-premises data center.

The Challenges

ESG asked IT professionals responsible for data storage technology to name the benefits their organization realized as the result of deploying on-premises flash storage technology. Improved performance, better TCO, and improved resource utilization were the most-cited responses (see Figure 1). In the same survey, respondents reported averages of 31% improved TCO, OpEx savings of 28%, and CapEx savings of 31% thanks to the use of flash storage. However, the overarching issues that drive data storage concerns remain relatively unchanged—data growth is accelerating, and the resulting infrastructure required to store and protect that data is costly and complex.¹

While flash leads to notable TCO improvements and cost savings in virtualized environments, for other use cases—video storage, file sharing and backup and archive, for example—a hybrid configuration might be a better fit. A storage solution should be flexible enough to support all-flash and hybrid configurations to offer the right configuration for each use case.

Figure 1. Benefits Realized as the Result of Deploying Flash Storage Technology



Source: Enterprise Strategy Group

¹ Source: ESG Master Survey Results, [2019 Data Storage Trends](#), November 2019.

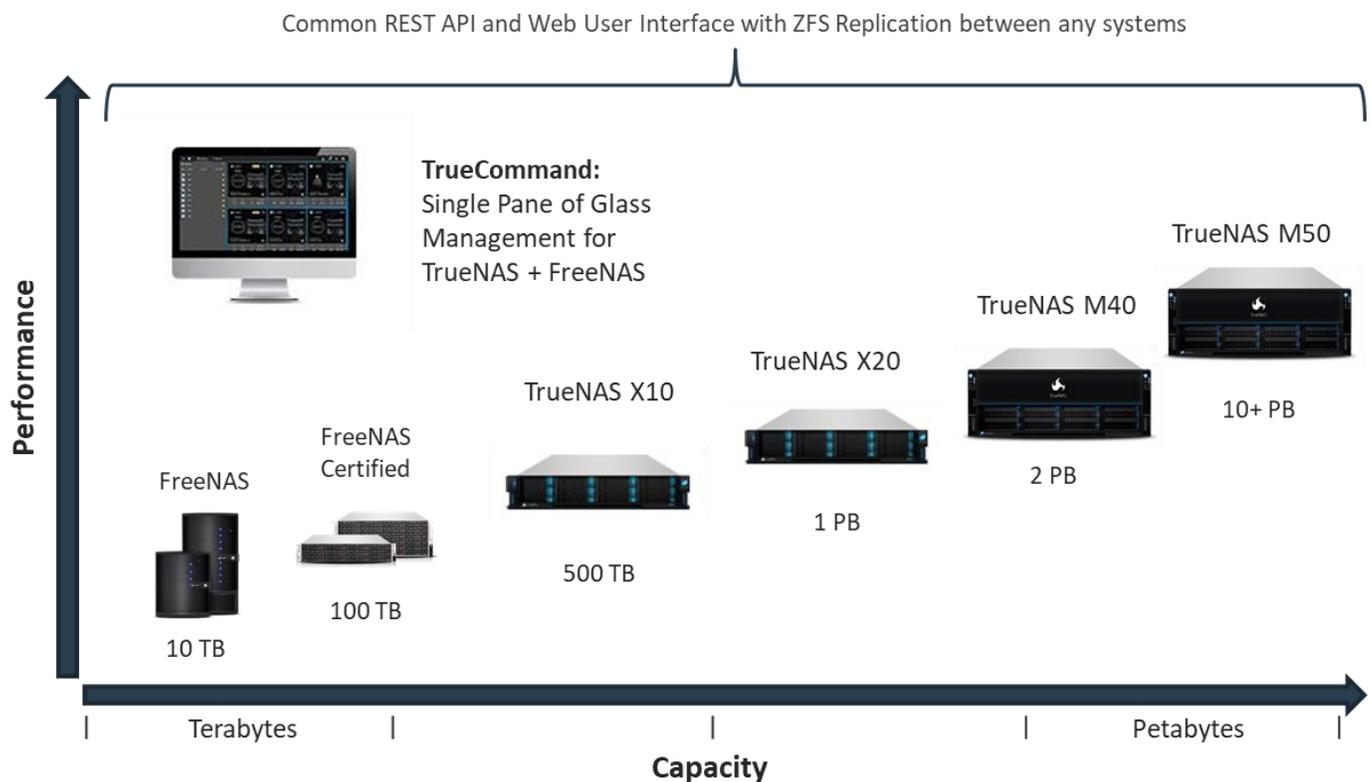
Organizations must provide a high-quality, predictable, and productive computing environment for an ever-growing number of internal users and external customers. In addition, enterprise application environments have become increasingly unpredictable as their underlying IT infrastructure grows in complexity and size. Business-critical application performance is sensitive to storage performance and latency, and highly dependent on the resilience of the enterprise IT environment.

The ability to consolidate mixed workloads and functions onto a single storage system has proven to provide significant TCO benefits if an organization is not meeting performance, reliability, and operational requirements.

The Solution: iXsystems TrueNAS

iXsystems has been providing storage and servers driven by open source technology, and is the developer of FreeNAS software-defined storage (SDS), an open-source operating system that users can install on virtually any hardware platform to share data over a network. With this deep expertise in hardware, software, and storage, iXsystems developed TrueNAS, an open storage platform designed to offer reliability and performance with strong cost efficiency.

Figure 2. The TrueNAS Scalable Product Family



Source: Enterprise Strategy Group

TrueNAS storage systems provide many storage features and functions that are must-haves in any enterprise.

Unified storage access over NAS (SMB and NFS), SAN (iSCSI and Fibre Channel), and S3-compatible APIs for object storage.

Self-correcting OpenZFS file system. When a read returns bad data, OpenZFS repairs it. Additional file system highlights:

- Thin and thick provisioning with quotas
- Copy-on-write file system ensures data is never overwritten and snapshots are efficient
- Intelligent in-line compression
- Encryption at rest with SED support and no impact to performance
- Two-tier caching: flash- or NVDIMM-based ZIL (ZFS Intent Log) for writes and DRAM-based ARC (adaptive replacement cache) and flash-based L2ARC (layer 2 adaptive replacement cache) for reads

- No limits on pool/share capacity or snapshots
- Integrated ZFS replication and Cloud Sync to AWS, Azure, Backblaze, and others

Scalability from single-digit terabytes to over 10 petabytes in a single system. Systems can be configured as all-flash for maximum performance and low latency and/or hybrid systems for lower cost per TB, with the same software and functionality available for each.

TrueCommand is iXsystems’ single pane of glass management platform designed to be easy to deploy and use; provide tools to simplify the scaling of data; and drive management and administration of iXsystems NAS platforms with role-based access control, collaborative management, and unified alerting and reporting.

ESG Tested

ESG tested TrueNAS to validate performance and functionality running multiple demanding workloads, including desktop virtualization, electronic design automation (EDA), and media and entertainment.

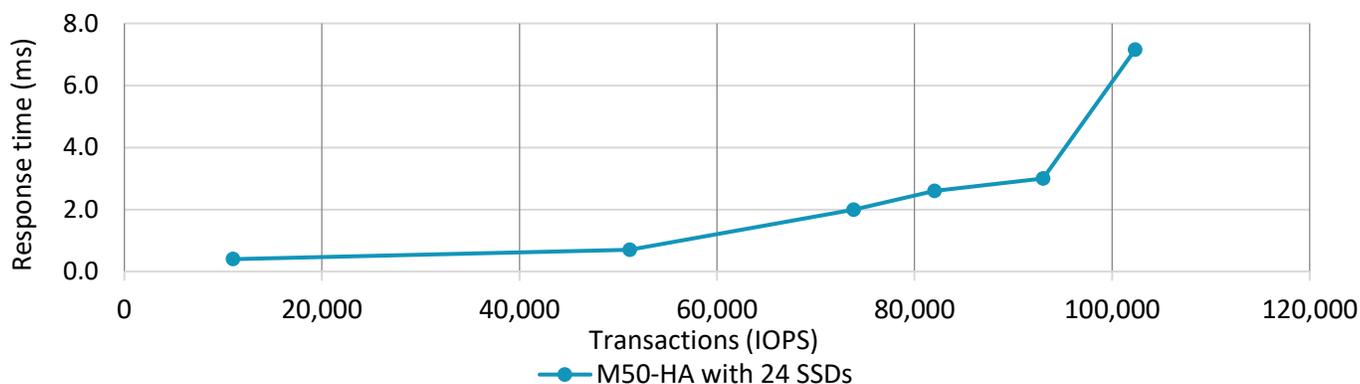
Desktop Virtualization

First, ESG examined TrueNAS in a desktop virtualization scenario servicing 5,000 power users with a very heavy I/O profile, using the Vdbench utility. The workload was based on the SPEC SFS 2014 VDI specification, with known NFS datastore artifacts removed and simplified I/O size ranges to make it work with Vdbench. This workload was designed to be as realistic as is practical for a synthetic workload. The majority of read and write I/O is random, with a single distribution spanning the entire LBA range, and distribution between hot and cold LBAs for all LUNs. Sequential I/O is done in batches, with new random offsets every 100 sequential I/Os to avoid artificially increasing cache hits. The overall read/write ratio was 27% reads and 73% writes with data compressibility set to 60%. Prior to the test runs, the LUNs were primed with incompressible data.

The configuration under test included a single TrueNAS M50-HA system with 24x 1.92TB SSDs in a Mirrored ZFS pool, and 12 iXsystems iXR31X12-HMC servers running CentOS Linux for control and load generation. Two Juniper QFX5200-32C 100Gbps switches provided connectivity. TrueNAS connected via two 100Gbps Ethernet connections (one from each controller), while all servers connected via a single 10Gbps Ethernet connection. Each server connected to its own iSCSI LUN via the Linux iSCSI initiator. Compression (LZ4) is enabled on each LUN/dataset by default, and it was used for these tests.

As seen in Figure 3, The TrueNAS M50 populated with only 24 SSDs exceeded 100,000 IOPS—with mixed I/O sizes averaging 24KB—at just over 7 ms response time. It’s important to note that testing measured performance of the system servicing a specific number of virtual desktop power users; this does not represent the maximum performance the platform can offer. ESG has observed much higher performance results for the TrueNAS M50 using synthetic workloads. The system serviced more than 850,000 IOPS using a 4KB, 100% random read workload and 9GB/s with 1MB, 100% sequential reads.

Figure 3. TrueNAS Virtualization Workload Response Time Curve



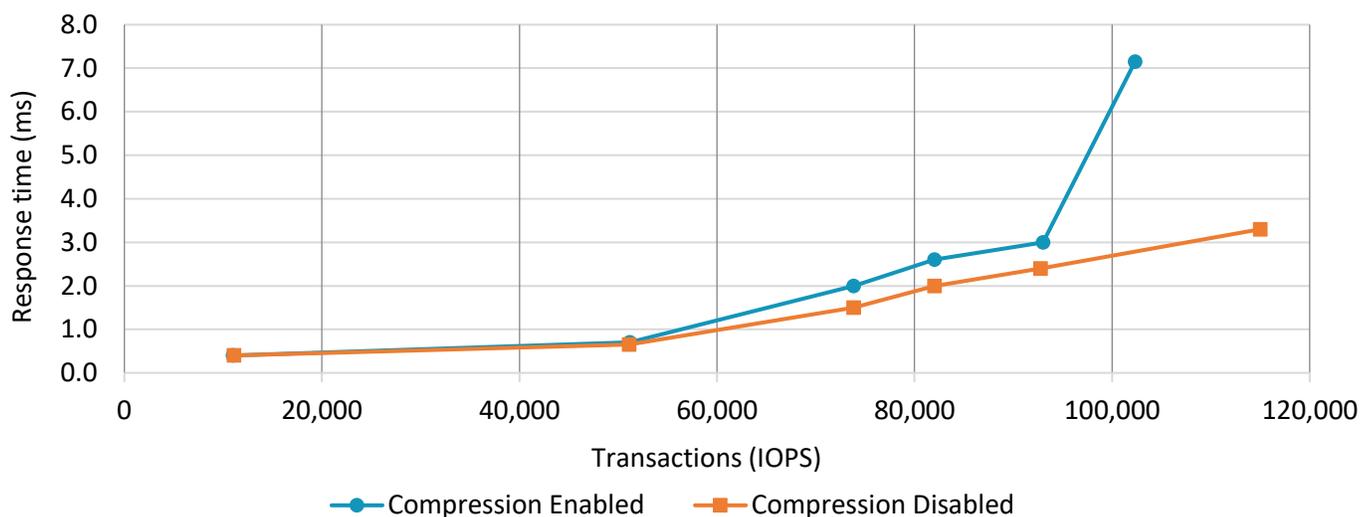
Source: Enterprise Strategy Group

Response time stayed under 1 ms to nearly 60,000 IOPS. Impressively, the M50 sustained more than 2,600 MiB/s of throughput with an average I/O size of 24 KB.

We also observed a significant reduction in disk IOPS and throughput as compared with application IOPS at higher workloads. At 100% load, disk IOPS were just 60% of application IOPS, and back-end disk throughput was just 70% of application throughput. This is also impressive considering that the workload is 73% writes and in a mirrored pool. Looking at the platform, we saw that metadata read requests were serviced almost completely from cache, which is a good thing since metadata responsiveness directly impacts file system performance. While we did not break out IOPS reduction into reads vs. writes, we did observe an increase in disk write size vs. application writes. It's reasonable to expect that this translates to write IOPS reduction, which would have a positive impact on SSD endurance.

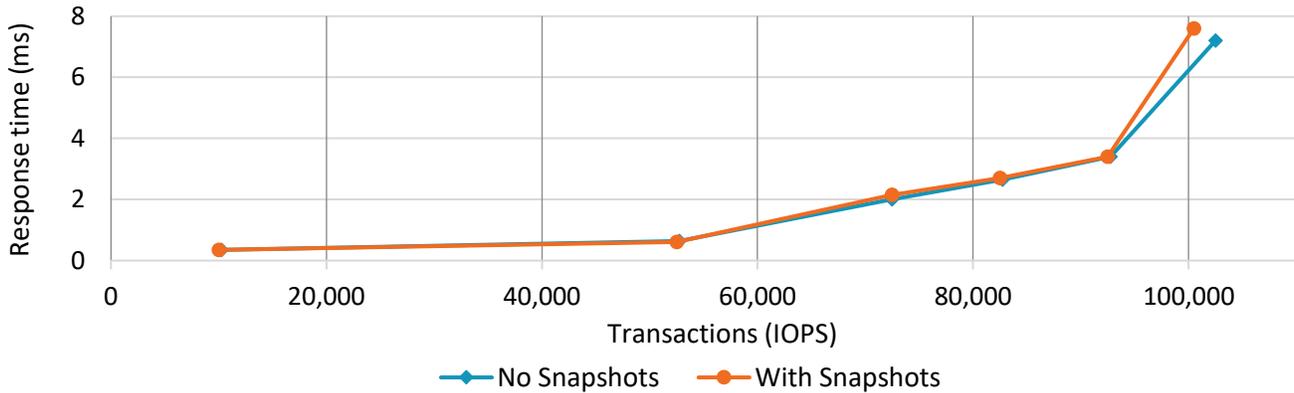
Next, we looked at the performance impact of compression. As seen in Figure 4, compression contributes to a slight response time increase through medium-high workloads, but compression is also contributing to a 33% capacity savings in our testing.

Figure 4. Performance Impact of Compression



Source: Enterprise Strategy Group

ESG repeated this test to examine the performance impact of snapshots. Because ZFS uses copy-on-write, we expected minimal performance impact. iXsystems explains that the primary observed impact is from increased space usage in the pool to retain overwritten/deleted data. Figure 5 shows the response time curve for the same virtualization test with snapshots enabled and disabled. For this test, we created a recursive snapshot of the parent dataset to all ZVOLs, which generated 13 snapshots every five minutes. Each load point ran for 10 minutes.

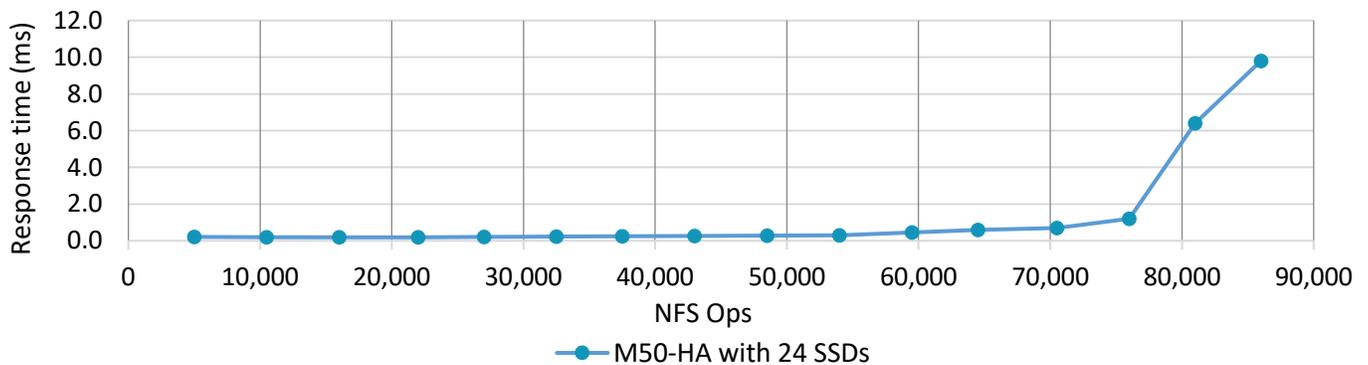
Figure 5. Performance Impact of Snapshots


Source: Enterprise Strategy Group

Impact in this scenario was, for all intents and purposes, nonexistent until the system hit 100% load, at which point the impact was minimal. We repeated the test, running at 50%, then 80% load for 55 minutes with the recursive snapshot task left running, with a new snap for every ZVOL every 5 minutes. We analyzed the data at one second intervals and confirmed zero impact to IOPS or response time over the course of the entire run.

Electronic Design Automation

Next, we tested file sharing performance using the NFS protocol. This phase of testing utilized a workload designed to measure the maximum throughput of a storage solution using file serving protocols to service an EDA workload comprising mixed file metadata and data operations with an average I/O size of 40 KiB. For this test, the file system record size was changed from the default of 128 KiB to 16 KiB. The load generating systems were populated with 32 GB of RAM to minimize client caching effects.

Figure 6. TrueNAS EDA Workload Response Time Curve


Source: Enterprise Strategy Group

The dual controller system serviced just over 70,000 NFS ops with sub 1ms response time.

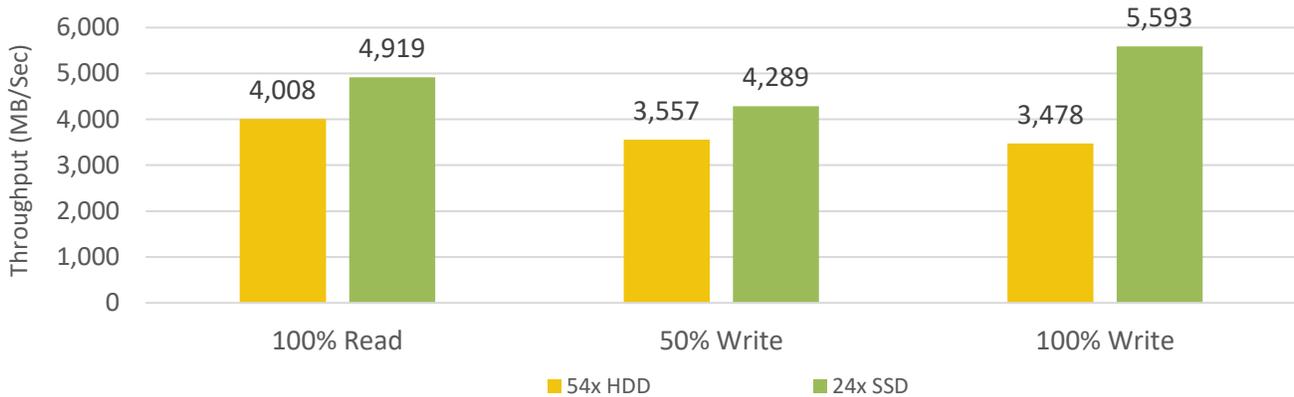
Sequential Workloads

Next, we tested sequential I/O, which is useful in media and entertainment, analytics (ingest), and seismic applications (for example)—basically, any time large volumes of bulk data need to be read or written. For this test, the same M50-HA was utilized with two configuration changes: the file system record size was changed from the default of 128KiB to 1 MiB, and the test was run on 24x 2 TB SSDs, then on 54x 2 TB SAS hard drives. The FIO benchmark (v. 3.13 for Windows) was used to generate the workloads, running on the same iXsystems iXR31X12-HMC servers, but this time running Windows 10 and

connecting to TrueNAS via SMB. We used a Dell workstation with dual Xeon E5-2660 CPUs and 128GB of RAM as our video editing workstation connecting to the single SMB share over the network with a 40GbE NIC.

We used a non-compressible dataset consisting of 12x 1.2 TB test files with the single client test using all 12 files sequentially and the multi-client test assigning a single file to each client.

Figure 7. Multiple Client Sequential Performance

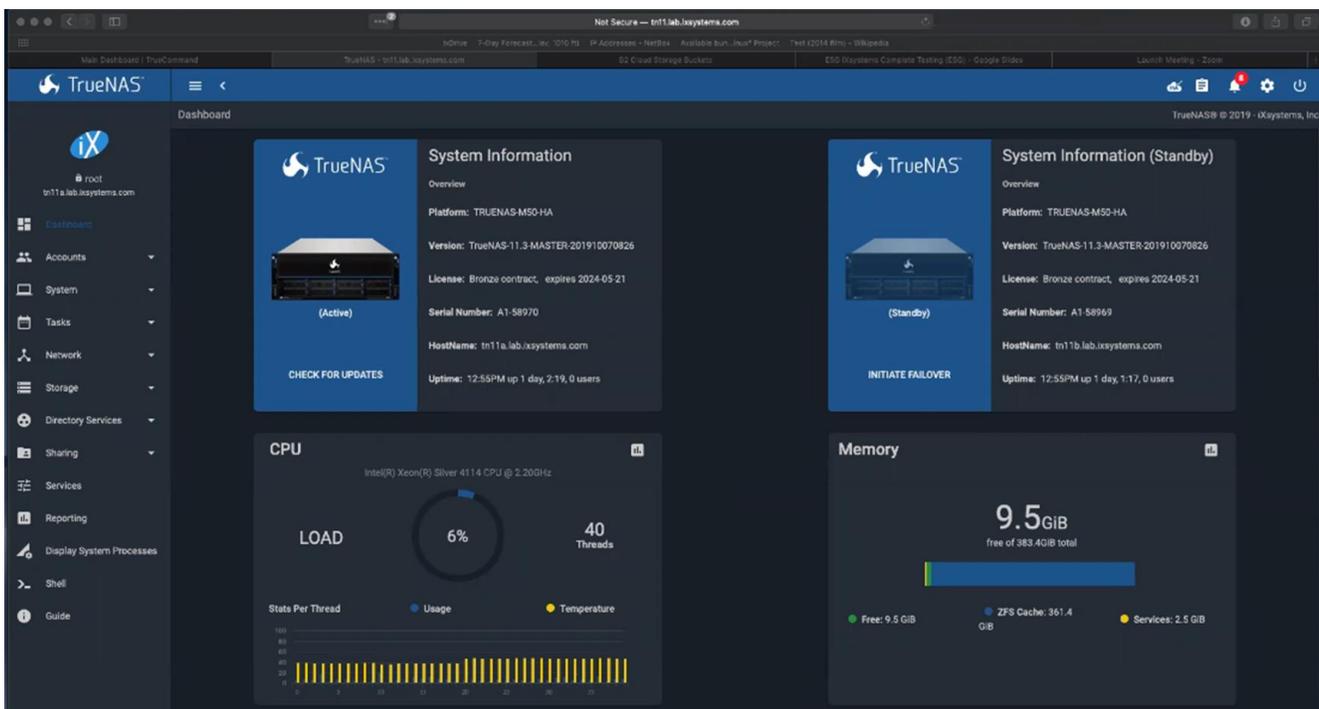


Source: Enterprise Strategy Group

Figure 7 shows the results of testing for multiple read mixes. The hard drive configuration provided consistent performance between 3,500 and 4,000 MiB/sec, and while the SSD configuration did exceed the hard drive performance, it would come at a cost premium. We also tested single-client, single-threaded performance using the hard drive configuration and confirmed that a single client with a single-threaded workload can achieve greater than 10Gb/sec of throughput.

Finally, we looked at the TrueNAS web UI (Figure 8) and TrueCommand management platform, shown in Figure 9.

Figure 8. TrueNAS Web User Interface

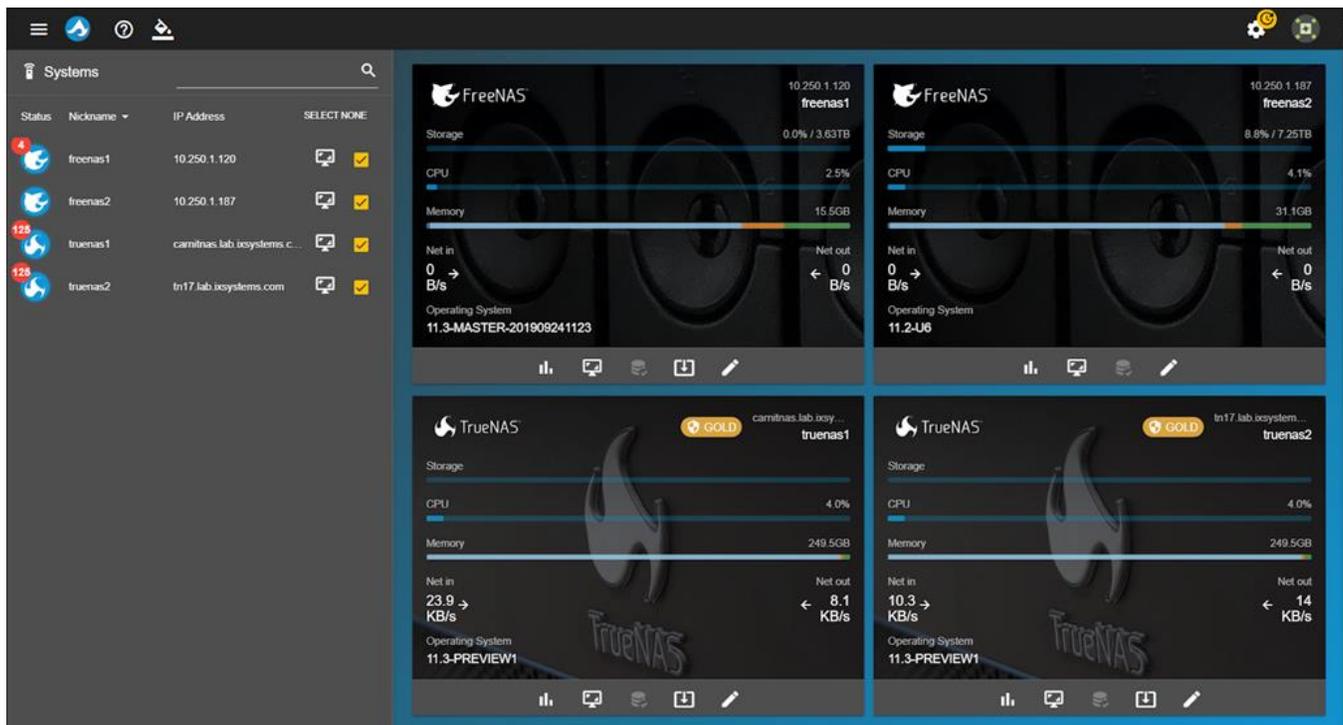


Source: Enterprise Strategy Group

We explored the TrueNAS web UI, performing multiple activities that an administrator would need to execute in a production environment, including:

- **Created and extended ZFS pools in mirrored and RAID-Z layouts.**
- **Configured cache.** The TrueNAS M50 uses NVDIMMs for write caching and NVMe U.2 dual ported drives for read caching. The M40 is similar but uses SAS SSDs for read caching. The TrueNAS X-Series (both X10 and X20) use SAS SSDs for both read and write caching.
- **Created shares, exports, and volumes, setting quotas, and creating a ZVOL**—a block LUN—which we could present over iSCSI or FibreChannel.
- **Configured snapshots, replication, and Cloud Sync.** Snapshots can be manual or scheduled, are pointer-based, operate nearly instantly when created, and don't occupy space. Administrators can clone a dataset or ZVOL snapshot to a fully functional dataset or ZVOL that is writeable, but only uses additional space for new data or overwrites. iXsystems does not set a limit on the number of snapshots that may be configured. Snapshots can be cloned to appear as writeable independent filesystems and/or volumes.
- **Walked through set up of replication and Cloud Sync.** Cloud Sync can be configured to push or pull data between TrueNAS and the cloud and can transfer the content using one of three tasks: *sync*, *copy*, or *move*. *Sync* copies any changes from the source system to the destination. *Copy* copies new files, and files deleted on the host are not deleted on the destination. *Move* copies files to the destination, then deletes them from the source after the transfer is complete. Cloud Sync can be used to enhance a disaster recovery (DR) scenario by extending replication to the cloud. As an example, TrueNAS "A" replicates data to TrueNAS "B" in another data center, then the data is backed up to Amazon S3 using Cloud Sync. If TrueNAS "B" goes away, TrueNAS "A" can recover data from Amazon S3. It's also important to note that iXsystems supports multiple cloud vendors for Cloud Sync including Microsoft Azure, Google Cloud, Backblaze, and many others.
- **We examined system and data health features**, including SMART tests, disk stats, disk scrubbing/re-silvering, and performance statistics and reporting. It's important to note that ZFS does not require file system checks—scheduled periodic SCRUB operations can walk the file system and correct any errors found.
- **TrueNAS provides VMware integration** via a vCenter plugin and coordinates with VMware snapshots. Coordinated snapshots between ESXi and TrueNAS provide a crash-consistent copy of VMs. TrueNAS instructs ESXi to pause the VM and take a snapshot. TrueNAS then maps the VM to the LUN or LUNs it occupies and takes a snapshot of them. The final step is to tell ESXi to discard the snapshot and un-pause the VM. TrueNAS keeps metadata that tells the user which VM is contained in the TrueNAS snapshot. Organizations can use replication technology to send snapshots to other systems or to the cloud, where it becomes a DR copy.

Finally, we looked at the TrueCommand management platform, shown in Figure 9. The TrueCommand dashboard provides instant access to high level system status. With TrueCommand, we could see the status of multiple systems and connect to the web UI of the specific TrueNAS under test through Single Sign On (SSO). Role Based Access Control (RBAC) determines which users can monitor and change specific systems.

Figure 9. TrueCommand Management Platform

Source: Enterprise Strategy Group

TrueCommand also provides two years of statistics collection, plus reporting and alert management. Any configuration changes to a TrueNAS system are logged for auditing and problem resolution.

Why This Matters

The complexity and performance challenges associated with a modern storage environment can lead to overbuying storage in order to achieve acceptable levels of performance. After all, slow storage can result in lost sales, lost customer goodwill, lost productivity, and lost competitiveness, and modern applications demand increasing levels of storage performance.

ESG has determined that iXsystems' TrueNAS delivers consistently impressive performance, and that it does so in a cost-efficient package. The solution accomplishes this through a combination of deep hardware, software, and storage expertise that enabled the design of a dual-controller architecture that is simpler and less expensive than competitive entries that leverage multiple controllers. The OpenZFS file system provides enterprise-class features developed over two decades and is designed to scale.

Business process improvement continues to be a high IT priority. One important area where businesses can improve process is system monitoring and management. ESG validated that the TrueNAS TrueCommand management platform makes monitoring and management of all TrueNAS and FreeNAS devices easy. Administrators can divide it up and assign pools of storage to project administrators, unburdening themselves of the day-to-day requirements of storage management. TrueCommand's ability to manage all of an enterprise's iXsystems NAS platforms in a single pane of glass simplifies overall management and allows administrators to focus on other issues.

The Bigger Truth

Explosive data growth and budget constraints are driving IT managers to rethink the way they have been doing storage. Legacy storage systems that are complex and expensive are displaced in favor of cost-efficient storage systems that are easier to manage. The cost and complexity of legacy storage system refreshes and forklift upgrades are at odds with corporate level mandates of reducing costs and increasing the return on investment.

The consistent performance of the iXsystems TrueNAS open storage platform enables organizations to deploy more capacity in a single storage system, which reduces overall costs. ESG has validated that the TrueNAS platform delivers impressive levels of cost-optimized performance. ESG was impressed with the performance, availability, and efficiency of the platform. Testing revealed that TrueNAS provides:

- High IOPS and NFS Ops with consistently low response times.
- High throughput utilizing SSDs and hard drives.
- Minimal to near-zero performance impact of compression and snapshots.
- A diverse and robust feature set, with role-based access and collaborative management, replication and Cloud Sync, unlimited snapshots, system and data health monitoring, and VMware integration.

ESG compared the cost of a TrueNAS solution with quotes from major storage vendors. Systems compared were similar in capacity, performance, and functionality. Factors taken into consideration included hardware, software, support, and maintenance costs. The combination of ZFS, open source, and a lean sales force provides efficiencies that enable iXsystems to offer high performance, feature-rich solutions at high value price points. We calculated a significant delta. Based on our testing and analysis, we found that iXsystems can typically offer up to double the performance at half the cost of their major competitors.

The results presented in this document are based on testing in a controlled environment. Due to the many variables in each production data center, it is important to perform planning and testing in your own environment to validate the viability and efficacy of any solution.

If you're considering an upgrade of a legacy primary or secondary storage system and you want to spend less and store more, ESG recommends that you consider the iXsystems TrueNAS storage platform with cloud extensibility, deep management analytics, and up to twice the performance at half the cost of the competition.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.