

TrueNAS SCALE

- 1: [Developer's Notes](#)
- 2: [Introduction](#)
 - 2.1: [SCALE Hardware Guide](#)
 - 2.2: [Software Development Life Cycle](#)
- 3: [Getting Started with SCALE](#)
 - 3.1: [Installing SCALE](#)
 - 3.2: [Migrating from TrueNAS CORE](#)
 - 3.3: [Post-install Configuration](#)
 - 3.4: [First Time Login](#)
 - 3.5: [Component Naming](#)
- 4: [Storage](#)
 - 4.1: [Pools](#)
 - 4.1.1: [Datasets](#)
 - 4.1.2: [Zvols](#)
 - 4.1.3: [Permissions](#)
 - 4.1.4: [Encryption](#)
 - 4.1.5: [Fusion Pools](#)
 - 4.2: [Snapshots](#)
 - 4.3: [Disks](#)
- 5: [Shares](#)
 - 5.1: [AFP Migration](#)
- 6: [Data Protection](#)
 - 6.1: [Scrub Tasks](#)
 - 6.2: [Cloud Sync Tasks](#)
 - 6.2.1: [How To Back Up Google Drive to TrueNAS SCALE](#)
 - 6.3: [Rsync Tasks](#)
 - 6.4: [Periodic Snapshot Tasks](#)
 - 6.5: [S.M.A.R.T. Tests](#)
 - 6.6: [Replication](#)
- 7: [Network](#)
- 8: [Credentials](#)
 - 8.1: [Local Users](#)
 - 8.2: [Local Groups](#)
 - 8.3: [Directory Services](#)
 - 8.4: [Backup Credentials](#)
 - 8.5: [Certificates](#)
 - 8.6: [2FA \(Two-Factor Authentication\)](#)
- 9: [Virtualization](#)
 - 9.1: [Creating and Managing VMs](#)
 - 9.2: [Accessing NAS From a VM](#)
- 10: [Apps](#)
 - 10.1: [Using SCALE Apps](#)
 - 10.2: [Using SCALE Catalogs](#)
 - 10.3: [Chia App](#)
 - 10.4: [Deploying TrueCommand on TrueNAS SCALE](#)
 - 10.5: [Using Docker on TrueNAS SCALE](#)
 - 10.6: [MinIO Clusters](#)
- 11: [Reporting](#)
- 12: [System Settings](#)
 - 12.1: [Update](#)
 - 12.2: [General Settings](#)
 - 12.3: [Advanced Settings](#)
 - 12.4: [Boot Environments](#)
 - 12.5: [Services](#)
 - 12.6: [Shell](#)
- 13: [SCALE API](#)
- 14: [Notices](#)
 - 14.1: [TrueNAS SCALE EULA](#)
 - 14.2: [ZFS Feature Flags Removal](#)
- 15: [SCALE Security Reports](#)
- 16: [User Recommendations](#)
 - 16.1: [Hardened Backup Repository for Veeam](#)



TrueNAS SCALE is the latest member of the TrueNAS family and provides Open Source HyperConverged Infrastructure (HCI) including Linux containers and VMs. TrueNAS SCALE includes the ability to cluster systems and provide scale-out storage with capacities of up to hundreds of Petabytes. Just like TrueNAS CORE, TrueNAS SCALE is designed to be the most secure and efficient solution to managing and sharing data over a network, from smaller home networks “scaled” up to massive business environments.

The Linux base of SCALE allows for a similar, but slightly different feature set that will appeal to an audience that is more familiar with Linux applications and workflows while TrueNAS CORE continues to provide the known and heavily tested performance and features from the FreeBSD operating system. SCALE is an acronym that represents the core features of the software:

Scaled-Out ZFS
Converged
Active-Active
Linux Containers
Easy to Manage

Unlike other HCI platforms, a user can get started with TrueNAS SCALE on a single node and incrementally scale up and scale out to over 100 storage nodes with many additional compute-only nodes. TrueNAS SCALE is true Disaggregated HCI, meaning storage and compute can be scaled independently. Each node can support Virtual Machines (with the KVM hypervisor) as well as Docker containers by using native Kubernetes.

Free to download and use, TrueNAS SCALE welcomes developers and testers to contribute to its Open Source development model. OpenZFS and Gluster combine to enable scale-out ZFS capabilities with excellent stability and very efficient compression and snapshots. Deploy a single hyperconverged node in a home/office or a cluster with hundreds of compute and storage nodes in a datacenter. With support for KVM VMs, Kubernetes, and Docker containers, it's easy to add applications to suit your every need.

1 - Developer's Notes

NOTE: This page is retired, please use the official documentation instead for TrueNAS SCALE 21.08 and later. Want to get involved in helping to collaborate on TrueNAS SCALE? Join our [Official Discord Server](#).

System Requirements

- Any x86_64 compatible (Intel or AMD) processor
- 8GB of RAM (More is better)
- 20GB Boot Device

Nightly Status

Nightly images for TrueNAS SCALE are built every 24 hours, at around 2AM Eastern (EDT/EST) time. Online updates are created every 2 hours and are available in the SCALE UI online updating page.

The Nightly ISO Image can be downloaded from: <https://download.truenas.com/truenas-scale-nightly/>

Users wanting to upgrade to a nightly image from 20.10 or 20.12 versions can do so by downloading the manual update file: <https://update.freenas.org/scale/TrueNAS-SCALE-Angelfish-Nightlies/TrueNAS-SCALE.update>

2 - Introduction

The Introduction topic includes our Software Development Life Cycle, our SCALE hardware guide, and our SCALE End User License Agreement

Ready to get started? Choose a topic or article from the left-side Navigation pane. Click the < symbol to expand the menu to show the topics under this section.

2.1 - SCALE Hardware Guide

- - [Minimum Hardware Requirements](#)
 - [Storage Considerations](#)
 - [Storage Media](#)
 - [Storage Solutions](#)
 - [Storage Device Sizing](#)
 - [Storage Device Burn-In](#)
 - [Storage Controllers](#)
 - [SAS Expanders](#)
 - [Storage Device Cooling](#)
 - [Memory, CPU, and Network Considerations](#)
 - [Memory Sizing](#)
 - [Error Correcting Code Memory](#)
 - [Central Processing Unit \(CPU\) Selection](#)
 - [Remote Management: IPMI](#)
 - [Power Supply Units](#)
 - [Uninterruptible Power Supplies](#)
 - [Ethernet Networking](#)
 - [High-Speed Interconnects](#)
 - [Virtualized TrueNAS CORE](#)

From repurposed systems to highly-custom builds, the fundamental freedom of TrueNAS is the ability to run it on almost any x86 computer.

Minimum Hardware Requirements

The recommended system requirements to install TrueNAS:

Processor	Memory	Boot Device	Storage
2-Core Intel 64-Bit or AMD x86_64 Processor	16 GiB Memory	16 GB SSD Boot Device	Two identically-sized devices for a single storage pool

The TrueNAS installer recommends 8 GB of RAM. TrueNAS installs, runs, operates jails, hosts SMB shares, and replicates TBs of data with less. iXsystems recommends the above for better performance and fewer issues.

You don't need an SSD boot device, but we discourage using a spinner or a USB stick for obvious reasons. We do not recommend installing TrueNAS on a single disk or striped pool unless you have a good reason to do so. You can install and run TrueNAS without any data device, but we strongly discourage it.

TrueNAS does not require two cores, as most halfway-modern 64-bit CPUs likely already have at least two.

For help building a system according to your unique performance, storage, and networking requirements, read on!

Storage Considerations

The heart of any storage system is the symbiotic pairing of its file system and physical storage devices. The ZFS file system in TrueNAS provides the [best available data protection of any file system at any cost](#) and makes very effective use of both spinning-disk and all-flash storage or a mix of the two. ZFS is prepared for the eventual failure of storage devices. It is highly configurable to achieve the perfect balance of redundancy and performance to meet any storage goal. A properly-configured TrueNAS system can tolerate the failure of multiple storage devices and even recreate its boot media with a copy of the [configuration file](#).

Storage Media

Choosing storage media is the first step in designing the storage system to meet immediate objectives and prepare for future capacity expansion.

Spinning Disks

Until the next scientific breakthrough in storage media, spinning hard disks are here to stay thanks to their balance of capacity and cost. The arrival of double-digit terabyte consumer and enterprise drives provides more choices to TrueNAS users than ever. TrueNAS Mini systems ship with Western Digital NAS and NL-SAS for good reason. Understanding the alternatives explains this decision.

SATA NAS Disks

Serial Advanced Technology Attachment (SATA) is still the de facto standard disk interface found in many desktop/laptop computers, servers, and some non-enterprise storage arrays. SATA disks first arrived offering double-digit gigabyte capacities and are now produced to meet many capacity, reliability, and performance goals. While consumer desktop SATA disks don't have the problematic overall reliability issues they once had, they are still not designed or warranted for continuous operation or use in RAID groups. Enterprise SATA disks address the always-on factor, vibration tolerance, and drive error handling required in storage systems. However, the price gap between desktop and enterprise SATA drives is vast enough that it forces users to push their consumer drives into 24/7 service to pursue cost savings.

Drive vendors, likely tired of honoring warranties for failed desktop drives used in incorrect applications, responded to this gap in the market by producing NAS drives. NAS drives achieved fame from the original Western Digital (WD) Red™ drives with CMR/PMR technology (now called WD Red Plus). Western Digital Designed the WD Red™ Plus NAS drives (non-SMR) for systems with up to 8 hard drives, the [WD Red™ Pro](#) for systems with up to 16 drives, and the [WD UltraStar™](#) for systems beyond 16 drives.

The iXsystems Community Forum regards WD drives as the preferred hard drives for TrueNAS builds due to their exceptional quality and reliability. All TrueNAS Minis ship with WD Red™ Plus drives unless requested otherwise.

Nearline SAS Disks

Nearline SAS (NL-SAS) disks are 7200 RPM enterprise SATA disks with the industry-standard SAS interface found in most enterprise storage systems. SAS stands for **Serial Attached SCSI**, with the traditional SCSI disk interface in serial form. SAS systems, designed for data center storage applications, have accurate, verbose error handling, predictable failure behavior, reliable hot swapping, and the added feature of multipath support. Multipath access means that each drive has two interfaces and can connect to two storage controllers or one controller over two cables. This redundancy protects against cable, controller card, or complete system failure in the case of the TrueNAS high-availability architecture in which each controller is an independent server that accesses the same set of NL-SAS drives. NL-SAS drives are also robust enough to handle the rigors of systems with more than 16 disks. So, capacity-oriented TrueNAS systems ship with [Western Digital UltraStar](#) NL-SAS disks thanks to the all-around perfect balance of capacity, reliability, performance, and flexibility that NL-SAS drives offer.

SAS Disks

Enterprise SAS disks, built for the maximum performance and reliability that a spinning platter can provide, are the traditional heavy-lifters of the enterprise storage industry. SAS disk capacities are low compared to NL-SAS or NAS drives due to the speed at which the platters spin, reaching as high as 15,000 RPMs. While SAS drives may sound like the ultimate answer for high-performance storage, many consumer and enterprise flash-based options have come onto the market and significantly reduced the competitiveness of SAS drives. For example, enterprise SAS drives discontinued from the TrueNAS product lines were almost completely replaced by flash drives (SSDs or NVMe) in 2016 due to their superior performance/cost ratio.

SATA & SAS Flash Storage SSDs

Flash storage technology has progressed significantly in recent years, leading to a revolution in mobile devices and the rise of flash storage in general-purpose PCs and servers. Unlike hard disks, flash storage is not sensitive to vibration and can be much faster with comparable reliability. Flash storage remains more expensive per gigabyte, but is becoming more common in TrueNAS systems as the price gap narrows.

The shortest path for introducing flash storage into the mainstream market was for vendors to use standard SATA/SAS hard disk interfaces and form factors that emulate standard hard disks but without moving parts. For this reason, flash storage Solid State Disks (SSDs) have SATA interfaces and are the size of 2.5" laptop hard disks, allowing them to be drop-in replacements for traditional hard disks. Flash storage SSDs can replace HDDs for primary storage on a TrueNAS system, resulting in a faster, though either a smaller or more expensive storage solution. If you plan to go all-flash, buy the highest-quality flash storage SSDs your budget allows with a focus on power, safety, and write endurance that matches your expected write workload.

NVMe

While SSDs pretending to be HDDs made sense for rapid adoption, the Non-Volatile Memory Express (NVMe) standard is a native flash protocol that takes full

advantage of flash storage's non-linear, parallel nature.

The main advantage of NVMe is generally its low-latency performance, and it's becoming a mainstream option for boot and other tasks. At first, NVMe was limited to expansion-card form factors such as PCIe and M.2. The new U.2 interface offers a universal solution that includes the 2.5" drive form factor and an externally accessible (but generally not hot-swappable) NVMe interface.

Note: NVMe devices can run quite hot and may need dedicated heat sinks.

USB Hard Disks

Avoid using USB-connected hard disks for primary storage with TrueNAS. You can use USB Hard Disks for very basic backups in a pinch. While TrueNAS does not automate this process, you can connect a USB HDD, replicate at the command line, and then take it off-site for safekeeping.

Warning: USB-connected media (including SSDs) may report their serial numbers inaccurately, making them indistinguishable from each other.

These storage device media arrange together to create powerful storage solutions.

Storage Solutions

Hybrid Storage & Flash Cache (SLOG/ZIL/L2ARC)

With hard disks providing double-digit terabyte capacities and flash-based options providing even higher performance, a best of both worlds option is available. With TrueNAS and OpenZFS, you can merge both flash and disk to create hybrid storage that makes the most of both storage types. Hybrid setups use high-capacity spinning disks to store data while DRAM and flash perform hyper-fast read and write caching. The technologies work together with a flash-based separate write log (SLOG). Think of it as a write cache keeping the ZFS-intent log (ZIL) used to speed up writes. On the read side, flash is a level two adaptive replacement (read) cache (L2ARC) to keep the hottest data sets on the faster flash media. Workloads with synchronous writes such as NFS and databases benefit from SLOG devices, while workloads with frequently-accessed data might benefit from an L2ARC device. An L2ARC device is not always the best choice because the level one ARC in RAM [always provide a faster cache](#), and the L2ARC table uses some RAM.

SLOG devices don't need to be large, since they only need to service five seconds of data writes delivered by the network or a local application. A high-endurance, low-latency device between 8 GB and 32 GB in size is adequate for most modern networks, and you can strip or mirror several devices for either performance or redundancy. Pay attention to the device's published endurance claims since a SLOG acts as the funnel point for most of the writes made to the system.

SLOG devices also need power protection. The purpose of the ZFS intent log (ZIL), and thus the SLOG, is to keep sync writes safe during a crash or power failure. If the SLOG isn't power-protected and loses data after a power failure, it defeats the purpose of using a SLOG in the first place. Check the manufacturer's specifications to ensure the SLOG device is power-safe or has power loss/failure protection.

The most important quality to look for in an L2ARC device is random read performance. The device needs to support more IOPS than the primary storage media it caches. For example, using a single SSD as an L2ARC is ineffective in front of a pool of 40 SSDs, as the 40 SSDs can handle far more IOPS than the single L2ARC drive. As for capacity, 5x to 20x larger than RAM size is a good guideline. High-end TrueNAS systems can have NVMe-based L2ARC in double-digit terabyte sizes.

Keep in mind that for every data block in the L2ARC, the primary ARC needs an 88-byte entry. Poorly-designed systems can cause an unexpected fill-up in the ARC and reduce performance in a p. For example, a 480 GB L2ARC filled with 4KiB blocks needs more than 10GiB of metadata storage in the primary ARC.

Self Encrypting Drives

TrueNAS supports two forms of data encryption at rest to achieve privacy and compliance objectives: [Native ZFS encryption](#) and [Self Encrypting Drives \(SEDs\)](#). SEDs do not experience the performance overhead introduced by software partition encryption but aren't as readily available as non-SED drives (and thus can cost a little more).

Boot Devices

Booting legacy FreeNAS systems from 8 GB or larger USB flash drives was once very popular. We recommend looking at other options since USB drive quality varies widely and modern TrueNAS versions perform increased drive writes to the boot pool. For this reason, all pre-built [TrueNAS Systems](#) ship with either M.2 drives or SATA DOMs.

SATA DOMs, or disk-on-modules, offer reliability close to that of consumer 2.5" SSDs with a smaller form factor that mounts to an internal SATA port and doesn't use a drive bay. Because SATA DOMs and motherboards with m.2 slots are not as common as the other storage devices mentioned here, users often boot TrueNAS systems from 2.5" SSDs and HDDs (often mirrored for added redundancy). The recommended size for the TrueNAS boot volume is 8 GB, but using 16 or 32 GB (or a 120 GB 2.5" SATA SSD) provides room for more boot environments.

Hot Swapability

TrueNAS systems come in all shapes and sizes. Many users want to have external access to all storage devices for efficient replacement if issues occur. Most hot-swap drive bays need a proprietary drive tray into which you install each drive. These bay and tray combinations often include convenient features like activity and identification lights to visualize activity and illuminate a failed drive with sesutil(8) (<https://www.freebsd.org/cgi/man.cgi?query=sesutil&sektion=8>) for CORE, (https://manpages.debian.org/testing/sg3-utils/sg3_utils.8.en.html) for SCALE). TrueNAS Mini systems ship with four or more hot-swap bays. TrueNAS R-Series systems can support dozens of drives in their head units and external expansion shelves. Pre-owned or repurposed hardware is popular among TrueNAS users.

Pay attention to the maximum performance offered by the hot-swap backplanes of a given system. Aim for at least 6 Gbps SATA III support. Note that hot-swapping PCIe NVMe devices is not currently supported.

Storage Device Sizing

[Zpool layout](#) (the organization of LUNs and volumes, in TrueNAS/ZFS parlance) is outside of the scope of this guide. The availability of double-digit terabyte drives raises a question TrueNAS users now have the luxury of asking: How many drives should I use to achieve my desired capacity? You can mirror two 16TB drives to achieve 16TB of available capacity, but that doesn't mean you should. Mirroring two large drives offers the advantage of redundancy and balancing reads between the two devices, which could lower power draw, but little else. The write performance of two large drives, at most, is that of a single drive. By contrast, an array of eight 4TB drives offers a wide range of configurations to optimize performance and redundancy at a lower cost. If configured as striped mirrors, eight drives could yield four times greater write performance with a similar total capacity. You might also consider adding a hot-spare drive with any zpool configuration, which lets the zpool automatically rebuild itself if one of its primary drives fails.

Storage Device Burn-In

Spinning disk hard drives have moving parts that are highly sensitive to shock and vibration and wear out with use. Consider pre-flighting every storage device before putting it into production, paying attention to:

- Start a long HDD self-test (`smartctl -t long /dev/`), and after the test completes (could take 12+ hrs)
- Check the results (`smartctl -a /dev/`)
- Check pending sector reallocations (`smartctl -a /dev/ | grep Current_Pending_Sector`)
- Check reallocated sector count (`smartctl -a /dev/ | grep Reallocated_Sector_Ct`)
- Check the UDMA CRC errors (`smartctl -a /dev/ | grep UDMA_CRC_Error_Count`)
- Check HDD and SSD write latency consistency (`diskinfo -wS`) *Unformatted drives only!*
- Check HDD and SSD hours (`smartctl -a /dev/ | grep Power_On_Hours`)
- Check NVMe percentage used (`nvmecontrol logpage -p 2 nvme0 | grep "Percentage used"`)

Take time to create a pool before deploying the system. Subject it to as close to a real-world workload as possible to reveal individual drive issues and help determine if an alternative pool layout is better suited to that workload. Be cautious of used drives as vendors may not be honest or informed about their age and health. Check the number of hours on all new drives using `smartctl(8)` to verify they aren't recertified. A drive vendor could also zero the hours of a drive during recertification, masking its true age. iXsystems tests all storage devices it sells for at least 48 hours before shipment.

Storage Controllers

The uncontested most popular storage controllers used with TrueNAS are the 6 and 12 Gbps (Gigabits per second, sometimes expressed as Gb/s) Broadcom (formerly Avago, formerly LSI) SAS host bus adapters (HBA). Controllers ship embedded on some motherboards but are generally PCIe cards with four or more internal or external SATA/SAS ports. The 6 Gbps LSI 9211 and its rebranded siblings that also use the LSI SAS2008 chip, such as the IBM M1015 and Dell H200, are legendary among TrueNAS users who build systems using parts from the second-hand market. Flash using the latest IT or Target Mode firmware to disable the optional RAID functionality found in the IR firmware on Broadcom controllers. For those with the budget, newer models like the Broadcom 9300/9400 series give 12 Gbps SAS capabilities and even NVMe to SAS translation abilities with the 9400 series. TrueNAS includes the `sas2flash`, `sas3flash`, and `storcli` commands to flash or perform re-flashing operations on 9200, 9300, and 9400 series cards.

Onboard SATA controllers are popular with smaller builds, but motherboard vendors are better at catering to the needs of NAS users by including more than the traditional four SATA interfaces. Be aware that many motherboards ship with a mix of 3 Gbps and 6 Gbps onboard SATA interfaces and that choosing the wrong one could impact performance. If a motherboard includes hardware RAID functionality, do not use or configure it, but note that disabling it in the BIOS might remove some SATA functionality depending on the motherboard. Most SATA compatibility-related issues are immediately apparent.

There are countless warnings against using hardware RAID cards with TrueNAS. ZFS and TrueNAS provide a built-in RAID that protects your data better than

any hardware RAID card. You can use a hardware RAID card if it's all you have, but there are limitations. First and most importantly, do not use their RAID facility if your hardware RAID card supports HBA mode, also known as passthrough or JBOD mode (there is one caveat in the bullets below). When used, it allows it to perform indistinguishably from a standard HBA. If your RAID card does not have this mode, you can configure a RAID0 for every single disk in your system. While not the ideal setup, it works in a pinch. If repurposing hardware RAID cards with TrueNAS, be aware that some hardware RAID cards:

- Could mask disk serial number and S.M.A.R.T. health information
- Could perform slower than their HBA equivalents
- Could cause data loss if using a write cache with a dead battery backup unit (BBU))

SAS Expanders

A direct-attached system, where every disk connects to an interface on the controller card, is optimal but not always possible. A SAS expander (a port multiplier or splitter) enables each SAS port on a controller card to service many disks. You find SAS expanders only on the drive backplane of servers or JBODs with more than twelve drive bays. For example, a [TrueNAS JBOD that eclipses 90 drives](#) in only four rack units of space wouldn't be possible without SAS expanders. Imagine how many eight-port HBAs you would need to access 90 drives without SAS expanders.

While SAS expanders, designed for SAS disks, can often support SATA disks via the SATA Tunneling Protocol or STP, we still prefer SAS disks for reasons mentioned in the NL-SAS section above (SATA disks function on a SAS-based backplane). Note that the opposite is not true: you can't use a SAS drive in a port designed for SATA drives.

Storage Device Cooling

A much-cited study floating around the Internet asserts that drive temperature has little impact on drive reliability. The study makes for a great headline or conversation starter, but carefully reading the report indicates that the drives were tested under optimal environmental conditions. The average temperature that a well-cooled spinning hard disk reaches in production is around 28 °C, and [one study](#) found that disks experience twice the number of failures for every 12 °C increase in temperature. Before adding drive cooling that often comes with added noise (especially on older systems), know that you risk throwing money away by running a server in a data center or closet without noticing that the internal cooling fans are set to their lowest setting. Pay close attention to drive temperature in any chassis that supports 16 or more drives, especially if they are exotic, high-density designs. Every chassis has certain areas that are warmer for whatever reason. Watch for fan failures and the tendency for some models of 8TB drives to run hotter than other drive capacities. In general, try to keep drive temperatures below the drive vendor's specification.

Memory, CPU, and Network Considerations

Memory Sizing

TrueNAS has higher memory requirements than many Network Attached Storage solutions for good reason: it shares [dynamic random-access memory](#) (DRAM or simply RAM) between sharing services, add-on plugins, jails, and virtual machines, and sophisticated read caching. RAM rarely goes unused on a TrueNAS system and enough RAM is key to maintaining peak performance. You should have at least 8 GB of RAM for basic TrueNAS operations with up to eight drives. Other use cases each have distinct RAM requirements:

- Add 1 GB for each drive added after eight to benefit most use cases.
- Add extra RAM (in general) if more clients will connect to the TrueNAS system. A 20 TB pool backing lots of high-performance VMs over iSCSI might need more RAM than a 200 TB pool storing archival data. If using iSCSI to back VMs, plan to use at least 16 GB of RAM for reasonable performance and 32 GB or more for optimal performance.
- Add 2 GB of RAM for directory services for the winbind internal cache.
- Add more RAM as required for plugins and jails as each has specific application RAM requirements.
- Add more RAM for virtual machines with a guest operating system and application RAM requirements.
- Add the suggested 5 GB per TB of storage for deduplication that depends on an in-RAM deduplication table.
- Add approximately 1 GB of RAM (conservative estimate) for every 50 GB of L2ARC in your pool. Attaching an L2ARC drive to a pool uses some RAM, too. ZFS needs metadata in ARC to know what data is in L2ARC.

Error Correcting Code Memory

Electrical or magnetic interference inside a computer system can cause a spontaneous flip of a single bit of RAM to the opposite state, resulting in a memory error. Memory errors can cause security vulnerabilities, crashes, transcription errors, lost transactions, and corrupted or lost data. So RAM, the temporary data storage location, is one of the most vital areas for preventing data loss.

Error-correcting code or ECC RAM detects and corrects in-memory bit errors as they occur. If errors are severe enough to be uncorrectable, ECC memory causes the system to hang (become unresponsive) rather than continue with errored bits. For ZFS and TrueNAS, this behavior virtually eliminates any chances that RAM errors pass to the drives to cause corruption of the ZFS pools or file errors.

The lengthy, Internet-wide debate on whether to use error-correcting code (ECC) system memory with OpenZFS and TrueNAS summarizes as:

- ECC RAM is *strongly* recommended as another data integrity defense

However:

- Some CPUs or motherboards support ECC RAM but not all
- Many TrueNAS systems operate every day without ECC RAM
- RAM of any type or grade can fail and cause data loss
- RAM is most likely to fail in the [first three months](#) so test all RAM before deployment.

Central Processing Unit (CPU) Selection

Choosing ECC RAM limits your CPU and motherboard options, but that can be a good thing. Intel® makes a point of limiting ECC RAM support to their lowest and highest-end CPUs, cutting out the mid-range i5 and i7 models.

Which CPU to choose can come down to a short list of factors:

- An underpowered CPU can create a performance bottleneck because of how OpenZFS does checksums, and compresses and (optional) encrypts data.
- A higher-frequency CPU with fewer cores usually performs best for SMB only workloads because of Samba, the lightly-threaded TrueNAS SMB daemon.
- A higher-core-count CPU is better suited for parallel encryption and virtualization.
- A CPU with AES-NI encryption acceleration support improves the speed of the file system and network encryption.
- A server-class CPU is recommended for its power and ECC memory support.
- A Xeon E5 CPU (or similar) is recommended for software-encrypted pools.
- An Intel Ivy Bridge CPU or later recommended for virtual machine use.

Watch for VT-d/AMD-Vi device virtualization support on the CPU and motherboard to pass PCIe devices to virtual machines. Be aware if a given CPU contains a GPU or requires an external one. Also, note that many server motherboards include a BMC chip with a built-in GPU. See below for more details on BMCs.

AMD CPUs are making a comeback thanks to the Ryzen and EPYC (Naples/Rome) lines. Support for these platforms is limited on FreeBSD and, by extension, TrueNAS CORE. However, Linux has significant support, and TrueNAS SCALE should work with AMD CPUs without issue.

Remote Management: IPMI

As a courtesy to further limit the motherboard choices, consider the Intelligent Platform Management Interface or IPMI (a.k.a. baseboard management controller, BMC, iLo, iDrac, and other names depending on the vendor) if you need:

- Remote power control and monitoring of remote systems
- Remote console shell access for configuration or data recovery
- Remote virtual media for TrueNAS installation or reinstallation

TrueNAS relies on its web-based user interface (UI), but you might occasionally need console access to make network configuration changes. TrueNAS administration and sharing default to a single network interface, which can be challenging when you need to upgrade features like LACP aggregated networking. The ideal solution is to have a dedicated subnet to access the TrueNAS web UI, but not all users have this luxury. The occasional visit to the hardware console is necessary for global configuration and even for system recovery. The latest TrueNAS Mini and R-Series systems ship with full-featured, HTML5-based IPMI support on a dedicated gigabit network interface.

Power Supply Units

The top criteria to consider for a power supply unit (or PSU) on a TrueNAS system are its:

- Power capacity (in watts) for the motherboard and number of drives it must support
- Reliability
- Efficiency rating

- Relative noise
- Optional redundancy to keep important systems running if one power supply fails

Select a PSU rated for the initial and a future load placed on it. Have a PSU with adequate power to migrate from a large-capacity chassis to a fully-populated chassis. Also, consider a hot-swappable redundant PSU to help guarantee uptime. Users on a budget can keep a cold spare PSU to limit their potential downtime to hours rather than days. A good, modern PSU is efficient and completely integrates into the IPMI management system to provide real-time fan, temperature, and load information.

Most power supplies carry a certified efficiency rating known as an [80 Plus](#) rating. The 80 plus rating indicates the power drawn from the wall is lost as heat, noise, and vibration, instead of doing useful work like powering your components. If a power supply needs to draw 600 watts from the wall to provide 500 watts of power to your components, it's operating at $500/600 = \sim 83\%$ efficiency. The other 100 watts get lost as heat, noise, and vibration. Power supplies with higher ratings are more efficient but also far more expensive. Do some return-on-investment calculations if you're unsure what efficiency to buy. For example, if an 80 Plus Platinum PSU costs \$50 more than the comparable 80 Plus Gold, it should save you at least \$10 per year on your power bill for that investment to pay off over five years. You can read more about 80 Plus ratings in [this post](#).

Uninterruptible Power Supplies

TrueNAS provides the ability to communicate with a battery-backed, uninterruptible power supply (UPS) over a traditional serial or USB connection to coordinate a graceful shutdown in the case of power loss. TrueNAS works well with APC brand UPSS, followed by CyberPower. Consider budgeting for a UPS with pure sine wave output. Some models of SSD can experience data corruption on power loss. If several SSDs experience simultaneous power loss, it could cause total pool failure, making a UPS a critical investment.

Ethernet Networking

The network in Network Attached Storage is as important as storage, but the topic reduces to a few key points:

- Simplicity - Simplicity is often the secret to reliability with network configurations.
- Individual interfaces - Faster individual interfaces such as 10/25/40/100GbE are preferable to aggregating slower interfaces.
- Interface support - Intel and Chelsio interfaces are the best-supported options.
- Packet fragmentation - Only consider a *jumbo frames* [MTU](#) with dedicated connections such as between servers or video editors and TrueNAS that are unlikely to experience packet fragmentation.
- LRO/LSO offload features - Interfaces with [LRO](#) and [LSO](#) offload features generally alleviates the need for jumbo frames and their use can result in lower CPU overhead.

High-Speed Interconnects

Higher band hardware is becoming more accessible as the hardware development pace increases and enterprises upgrade more quickly. Home labs can now deploy and use 40 GB and higher networking components. Home users are now discovering the same issues and problems with these higher speeds found by Enterprise customers.

iXsystems recommends using optical fiber over *direct attached copper* (DAC) cables for the high speed interconnects listed below:

- 10Gb NICs: SFP+ connectors
- 25Gb NICs: SFP28 connectors
- 40Gb NICs: QSFP+ connectors
- 100Gb NICs: QSFP28 connectors
- 200Gb NICs: QSFP56 connectors
- 400Gb NICs: QSFP-DD connectors

iXsystems also recommends using optical fiber for any transceiver form factors mentioned when using fiber channels. Direct attached copper (DAC) cables could create interoperability issues between the NIC, cable, and switch.

Virtualized TrueNAS CORE

Finally, the ultimate TrueNAS hardware question is whether to use actual hardware or choose a virtualization solution. TrueNAS developers [virtualize TrueNAS every day](#) as part of their work, and cloud services are popular among users of all sizes. TrueNAS's design has OpenZFS at its heart. The design from day one works with physical storage devices. It is aware of their strengths and compensates for their weaknesses. When the need arises to virtualize TrueNAS:

- Pass hardware disks or the entire storage controller to the TrueNAS VM if possible (requires VT-d/AMD-Vi support).
- Disable automatic scrub pools on virtualized storage such as VMFS, and never scrub a pool while also running storage repair tasks on another layer.
- Use a least three vdevs to provide adequate metadata redundancy, even with a striped pool.
- Provide one or more 8 GB or larger boot devices.
- Provide the TrueNAS VM with adequate RAM per its usual requirements.
- Consider jumbo frame networking if all devices support it.
- Understand that the guest tools in FreeBSD might lack features found in other guest operating systems.
- Enable MAC address spoofing on virtual interfaces and enable promiscuous mode to use VNET jail and plugins.

2.2 - Software Development Life Cycle

- [SDLC Application](#)
- [TrueNAS Quality Lifecycle](#)

The TrueNAS (and FreeNAS) Software Development Life Cycle (SDLC) is the process of planning, creating, testing, deploying, and maintaining TrueNAS releases.

There are five stages to the TrueNAS SDLC: requirement analysis, design and development, testing and evaluation, documentation, and maintenance.

Requirement Analysis

Determine the objectives, nature, and scope of future versions of the software. Requirement Analysis involves gathering feedback and interpreting customer needs and requirements, diagnosing existing problems, and weighing the pros and cons of potential solutions. The end result is a list of recommended improvements to be integrated into future versions of TrueNAS.

Design and Development

Required and planned changes are investigated in detail and development steps are determined. Proposed alterations are reviewed by peers for completeness, correctness, and proper coding style. TrueNAS developers then begin altering the software to include new features, resolve software bugs, or implement security improvements.

Testing and Evaluation

Code is integrated into the existing TrueNAS source tree, then built and tested by the Release Engineering (RE) department. RE verifies that all requirements and objectives are properly met and the updated software is reliable and fault-tolerant according to the determined requirements. If issues are found, code is reworked to meet the development requirements. Simultaneously, a security evaluation of the TrueNAS code is completed, with any discovered issues sent to the engineering team for resolution.

Documentation

The Validation and Documentation Team audits all development changes to the software and resolves any inconsistencies with the current software documentation. This is to verify that end user documentation is as accurate as possible. Any security notices, errata, or best practices are also drafted for inclusion on the [TrueNAS Security website](#).

Maintenance

The new release of TrueNAS is evaluated to determine further feature development, bug fixes, or security vulnerability patches. During this stage, security patches and software erratum are corrected, updated versions of existing branches are pushed, and feedback is solicited for future versions of the software.

SDLC Application

The TrueNAS SDLC applies to the latest two release branches. As new releases are created for TrueNAS, the oldest TrueNAS release branch is dropped out of the SDLC and labeled as End of Life (EoL). For example, TrueNAS/FreeNAS 11.3 and TrueNAS 12.0 were in active development under the SDLC in August 2020. In early 2021, TrueNAS Core/Enterprise 12.0 and 12.1 branches were in active development under the SDLC. These versions of the software are in active development and maintenance. We encourage users to actively keep their software updated to an active development version to continue to receive security patches and other software improvements.

TrueNAS Quality Lifecycle

TrueNAS releases follow a general adoption guideline for their lifetime. Starting with the NIGHTLY builds, each stage of a major release incorporates more testing cycles and bug fixes that represent a maturation of the release. With each version release stage, users are encouraged to install, upgrade, or otherwise begin using the major version, depending on the specific TrueNAS deployment and use case:

Release Stage	Completed QA Cycles	Typical Use-case	Description
NIGHTLY	0	Developers	Incomplete
ALPHA	1	Testers	Not much field testing
BETA	2	Enthusiasts	Major Feature Complete, but expect some bugs
RC	3	Home Users	Suitable for non-critical deployments
RELEASE	4	General Use	Suitable for less complex deployments
U1	5	Business Use	Suitable for more complex deployments
U2+	6+	Mission Critical	Suitable for critical uptime deployments

3 - Getting Started with SCALE



This section guides you through installing and accessing TrueNAS SCALE, storing, backing up, and sharing data, and expanding TrueNAS with different applications solutions.

Ready to get started? Choose a topic or article from the left-side Navigation pane. Click the < symbol to expand the menu to show the topics under this section.

For more detailed interface reference articles, configuration instructions, and tuning recommendations, see the remaining sections in this topic. Content sections are organized by order of appearance in the web interface.

3.1 - Installing SCALE

Once you have [downloaded](#) the .iso file, you can start installing TrueNAS SCALE!

ISO Verification

The iXsystems Security Team cryptographically signs TrueNAS .iso files so that users can verify the integrity of their downloaded file. This section demonstrates how to verify an .iso file using the [Pretty Good Privacy \(PGP\)](#) and [SHA256](#) methods.

PGP ISO Verification

You will need an OpenPGP encryption application for this method of ISO verification. There are many different free applications available, but the OpenPGP group provides a list of available software for different operating systems at <https://www.openpgp.org/software/>. The examples in this section show verifying the TrueNAS .iso using [gnupg2](#) in a command prompt, but [Gpg4win](#) is also a good option for Windows users.

To verify the .iso source, go to <https://www.truenas.com/download-tn-scale/>, expand the **Security** option, and click *PGP Signature* to download the Gnu Privacy Guard (.gpg) signature file. Open the [PGP Public key link](#) and note the address in your browser and **Search results for** string .

Use one of the OpenPGP encryption tools mentioned above to import the public key and verify the PGP signature.

Go to the .iso and .iso.gpg download location and import the public key using the keyserver address and search results string:

```
user@ubuntu /tmp> gpg --keyserver keys.gnupg.net --recv-keys 0xc8d62def767c1db0dff4e6ec358eaa9112cf7946
gpg: DBG: Using CREATE_BREAKAWAY_FROM_JOB flag
gpg: key 358EAA9112CF7946: public key "IX SecTeam <security-officer@ixsystems.com>" imported
gpg: DBG: Using CREATE_BREAKAWAY_FROM_JOB flag
gpg: Total number processed: 1
gpg:      imported: 1
user@ubuntu /tmp>
```

Use `gpg --verify` to compare the .iso and .iso.gpg files:

```
user@ubuntu /tmp> gpg --verify TrueNAS-SCALE-21.04-ALPHA.1.iso
gpg: Signature made Thu May 27 10:49:02 2021 EDT using RSA key ID 12CF7946
gpg: Good signature from "IX SecTeam <security-officer@ixsystems.com>"
gpg: WARNING: This key is not certified with a trusted signature!
gpg:          There is no indication that the signature belongs to the owner.
Primary key fingerprint: C8D6 2DEF 767C 1DB0 DFF4 E6EC 358E AA91 12CF 7946
user@ubuntu /tmp>
```

This response means the signature is correct but still untrusted. Go back to the browser page that has the **PGP Public key** open and manually confirm that the key was issued for IX SecTeam <security-officer@ixsystems.com> (iX Security Team) on October 15, 2019 and has been signed by an iXsystems account.

SHA256 Verification

The command to verify the checksum varies by operating system:

- BSD: `sha256 isofile`
- Linux: `sha256sum isofile`
- Mac: `shasum -a 256 isofile`
- Windows or Mac users can install additional utilities like [HashCalc](#) or [HashTab](#).

The value produced by running the command must match the value shown in the sha256.txt file. Different checksum values indicate a corrupted installer file that should not be used.

Choose the install type to see specific instructions:

Physical Hardware

Hardware Considerations

TrueNAS SCALE is very flexible and can run on any x86_64 compatible (Intel or AMD) processor. SCALE requires at least 8GB of RAM (more is better) and a 20GB Boot Device. If you're still researching what kind of hardware to use with SCALE, read over the very detailed [SCALE Developer's Notes](#).

Prepare the Install File

Physical hardware requires burning the TrueNAS SCALE installer to a device, typically a CD or removable USB device. This device is temporarily attached to the system to install TrueNAS SCALE to the system's permanent boot device.

To write the TrueNAS installer to a USB stick on Linux, plug the USB stick into the system and open a terminal.

Start by making sure the USB stick connection path is correct. There are many ways to do this in Linux, but a quick option is to enter `lsblk -po +vendor,model` and note the path to the USB stick. This shows in the **NAME** column of the `lsblk` output.

Next, use `dd` to write the installer to the USB stick.

Be very careful when using `dd`, as choosing the wrong `of=` device path can result in irretrievable data loss!

Enter `dd status=progress if=path/to/.iso of=path/to/USB` in the CLI. If this results in a "permission denied" error, use `sudo dd` with the same parameters and enter the administrator password.

Install Process

With the installer added to a device, you can now install TrueNAS SCALE onto the desired system. Insert the install media and reboot or boot the system. At the motherboard splash screen, use the hotkey defined by your motherboard manufacturer to boot into the motherboard UEFI/BIOS.

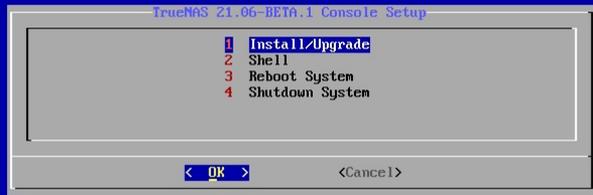
Choose to boot in UEFI mode or legacy CSM/BIOS mode. When installing TrueNAS, make the matching choice for the installation. For Intel chipsets manufactured in 2020 or later, UEFI is likely the only option.

If your system supports SecureBoot, you will need to either disable it or set it to "Other OS" to be able to boot the install media.

Select the install device as the boot drive, exit, and reboot the system. If the USB stick is not shown as a boot option, try a different USB slot. Which slots are available for boot differs by hardware.

After the system has booted into the installer, follow these steps.

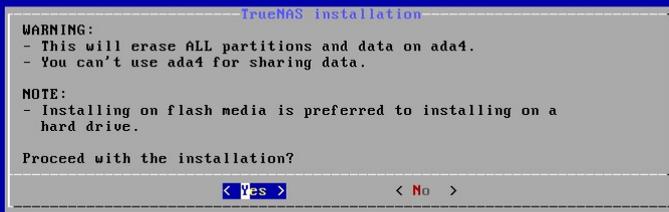
Select *Install/Upgrade*.



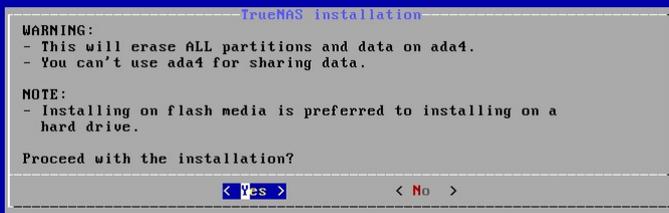
Select the desired install drive.



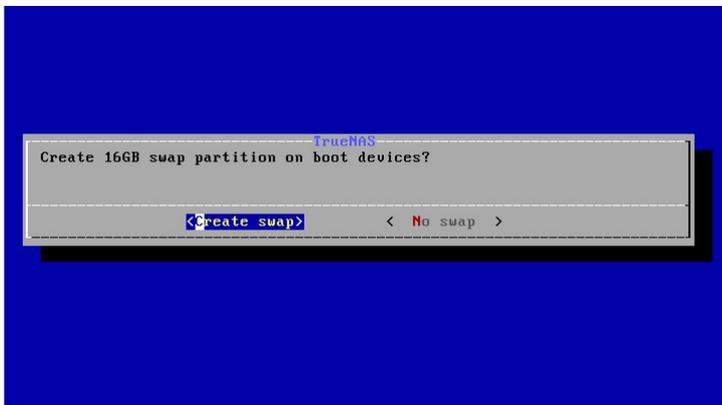
Select Yes.



Select *Fresh Install* to do a clean install of the downloaded version of TrueNAS SCALE. **This will erase the contents of the selected drive.!**



When the operating system device has enough additional space, you can choose to allocate some space for a swap partition to improve performance.



Enter a password for the root user to log in to the web interface.



After following the steps to install, reboot the system and remove the install media.

Troubleshooting

If the system does not boot into TrueNAS SCALE, there are several things that can be checked to resolve the situation:

- Check the system BIOS and see if there is an option to change the USB emulation from CD/DVD/floppy to hard drive. If it still will not boot, check to see if the card/drive is UDMA compliant.
- If the system BIOS does not support EFI with BIOS emulation, see if it has an option to boot using legacy BIOS mode.
- If the system starts to boot but hangs with this repeated error message: `run_interrupt_driven_hooks: still waiting after 60 seconds for xpt_config`, go into the system BIOS and look for an onboard device configuration for a 1394 Controller. If present, disable that device and try booting again.
- If the burned image fails to boot and the image was burned using a Windows system, wipe the USB stick before trying a second burn using a utility such as [Active@ KillDisk](#). Otherwise, the second burn attempt will fail as Windows does not understand the partition which was written from the image file. Be very careful to specify the correct USB stick when using a wipe utility!

Virtual Machine



Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaleangelfishyminstall.mp4>

Play Video

Because TrueNAS SCALE is built and provided as an .iso file, it works on all virtual machine solutions (VMware, VirtualBox, Citrix Hypervisor, etc). This section demonstrates installing with [VMware Workstation Player](#) on Windows.

Minimum Virtual Machine Settings

Regardless of virtual machine solution, use these minimum settings:

- RAM: at least 8192MB (8GB)
- DISKS: one virtual disk with at least 8GB for the operating system and boot environments and at least one additional virtual disk with at least 4GB to be used as data storage
- NETWORK: Use NAT, Bridged, or Host-only depending on your host network configuration.

Networking checks for VMware

When installing TrueNAS in a VMware VM, double check the virtual switch and VMware port group. Network connection errors for plugins or jails inside the TrueNAS VM can be caused by a misconfigured virtual switch or VMware port group. Make sure *MAC spoofing* and *promiscuous mode* are enabled on the switch first, and then the port group the VM is using.

Jail Networking

If you have installed TrueNAS in VMware, you will need functional networking to create a jail.

For the jail to have functional networking, you have to change the VMware settings to allow Promiscuous, MAC address changes, and Forged Transmits.

Setting	Description
Promiscuous Mode	When set to Accept , all traffic on the virtual switch level, objects defined within all portgroups can receive all incoming traffic on the vSwitch.
MAC Address Changes	When set to Accept , ESXi accepts requests to change the effective MAC address to a different address than the initial MAC address.
Forged Transmits	When set to Accept , ESXi does not compare source and effective MAC addresses.

Generic VM Creation Process

For most hypervisors, the procedure for creating a TrueNAS VM is the same:

1. Create a new Virtual Machine, as usual, taking note of the following settings.
2. The virtual hardware has a bootable CD/DVD device pointed to the TrueNAS SCALE installer image (this is usually an .iso).
3. The virtual network card is configured so it can be reached from your network. **bridged** mode is optimal as this treats the network card as if it is plugged into a simple switch on the existing network.
4. Some products require identifying the OS being installed on the VM. The ideal option is *Debian 11 64 bit*. If this is not available, try options like *Debian 11, Debian 64 bit, 64 bit OS, or Other*. **Do not choose a Windows, Mac or BSD related OS type.**
5. For VMware hypervisors, install in BIOS mode.
6. The VM has sufficient memory and disk space. TrueNAS needs at least 8 GB RAM and 20 GB disk space. Not all hypervisors allocate enough memory by default.
7. Boot the VM and install TrueNAS as usual.
8. When installation is complete, shut down the VM instead of rebooting, and disconnect the CD/DVD from the VM before rebooting the VM.
9. After rebooting into TrueNAS, install VM tools if applicable for your VM, and if they exist for Debian 11, or ensure they are loaded on boot.

Example installation for VMWare Player 15.5

Open VMware Player and click *Create a New Virtual Machine* to enter the New Virtual Machine Wizard.

1. Installer disk image file

Select the *Installer disk image file (.iso)* option, click *Browse...*, and upload the TrueNAS SCALE .iso downloaded earlier.

2. Name the Virtual Machine

In this step, the virtual machine name and location can be changed.

3. Specify Disk Capacity

Specify the maximum disk size for the initial disk. The default *20GB* is enough for TrueNAS. Next, select *Store virtual disk as a single file*.

4. Review Virtual Machine

Review the virtual machine configuration before proceeding. By default, VMware Player doesn't set enough RAM for the virtual machine. Click *Customize Hardware... > Memory*. Drag the slider up to 8GB and click *Ok*. If you wish to power on the machine after creation, select *Power on this virtual machine after creation*.

Add Virtual Disks for Storage

After the virtual machine has been created, select it from the virtual machine list and click *Edit virtual machine settings*. Click *Add...* and select *Hard Disk*. Select *SCSI* as the virtual disk type. Select *Create a new virtual disk*. Specify the maximum size of this additional virtual disk. This disk stores data in TrueNAS. If desired, allocate the disk space immediately by setting *Allocate all disk space now*. Select *Store virtual disk as single file*. Finally, name and chose a location for the new virtual disk.

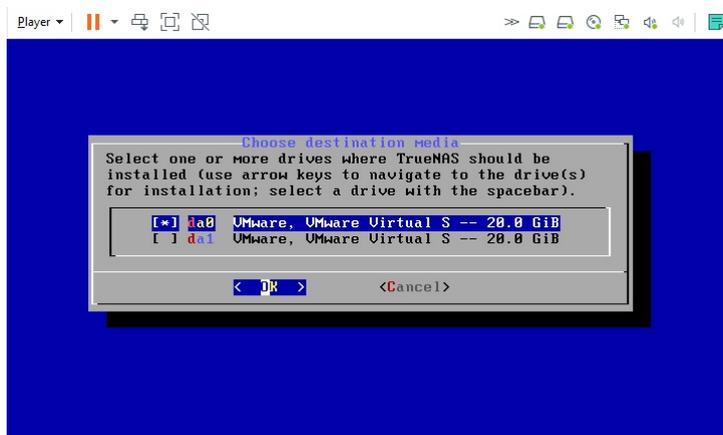
Repeat this process until enough disks are available for TrueNAS to create ideal storage pools This depends on your specific TrueNAS use case. See [Pool Creation](#) for descriptions of the various pool ("vdev") types and layouts

TrueNAS Installer

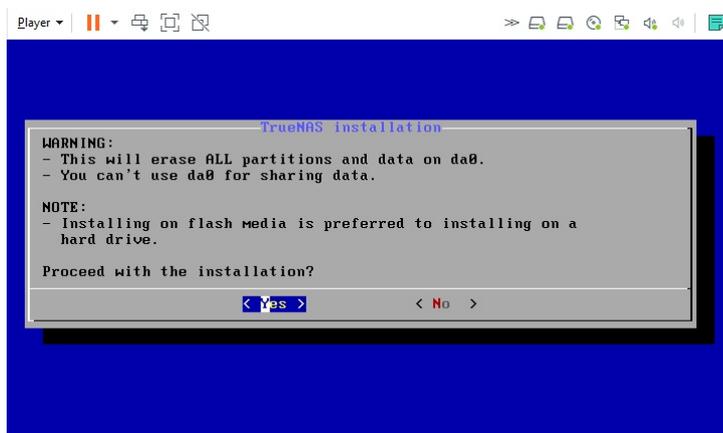
Select the virtual machine from the list and click *Play virtual machine*. The machine starts and boots into the TrueNAS installer. Select *Install/Upgrade*.



Select the desired disk for the boot environments.



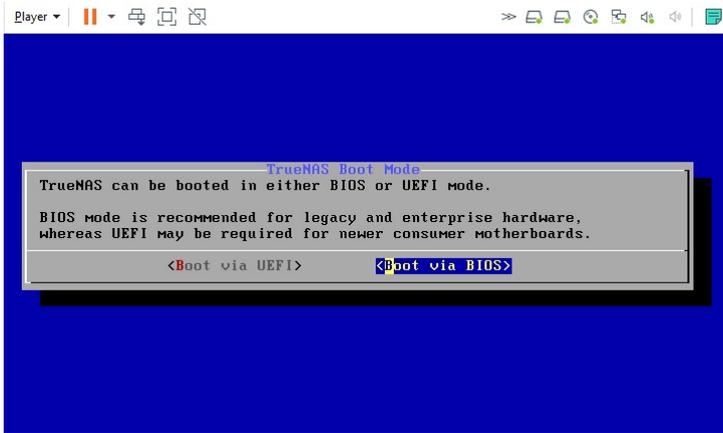
Select Yes. **This will erase all contents on the disk!**



Set a password for root login.



Select *Boot via BIOS*.



After the TrueNAS SCALE installation is complete, reboot the system. The [Console Setup Menu](#) displays when the system boots successfully.

3.2 - Migrating from TrueNAS CORE

- [Parallel SCALE CLI Commands](#)

Migrating TrueNAS from CORE to SCALE is a one-way operation. Attempting to activate or roll back to a CORE boot environment can break the system.

Migrating GELI-encrypted Pools to SCALE

TrueNAS SCALE is based on Linux, which does not support FreeBSD GELI encryption. If you have GELI-encrypted pools on your system that you plan to import into SCALE, you must migrate your data from the GELI pool to a non-GELI encrypted pool *before* migrating to SCALE.

ISO File

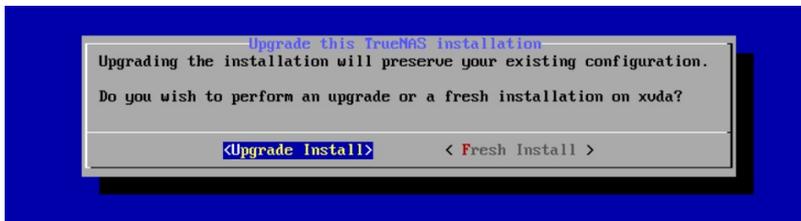
Start by saving the [SCALE ISO file](#) to a USB drive (see the **Physical Hardware** tab in [Installing SCALE](#)). Plug the USB drive into the CORE system that you want to sidegrade and boot or reboot the system.

At the motherboard splash screen, use the hotkey defined by your motherboard manufacturer to select a boot device, then select the USB drive with the SCALE .iso.

When the SCALE console setup screen appears, select **Install/Upgrade**.

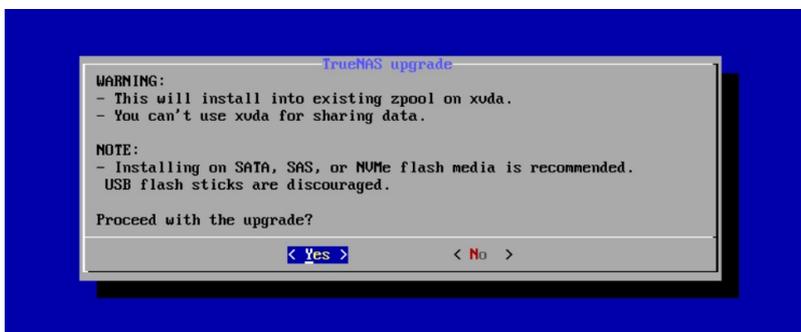


The installer asks if you want to preserve your existing configuration or start with a fresh installation. We recommend selecting *Upgrade Install* when migrating from CORE to SCALE to keep your configuration data. Then select **Install in new boot environment**.



Although TrueNAS attempts to keep most of your CORE configuration data when upgrading to SCALE, some CORE-specific items do not transfer. GELI Encrypted pools, NIS data, metadata, jails, tunables, and boot environments do not migrate from CORE to SCALE. AFP shares also do not transfer, but can be migrated into an SMB share with AFP compatibility enabled. Init/shutdown scripts transfer, but can break and should be reviewed before use. The CORE netcli utility is also swapped for a new CLI utility that is used for the Console Setup Menu and other commands issued in a CLI.

After choosing to install in new boot environment, the installer warns that SCALE installs into the boot pool previously used for CORE. Select **Yes**.



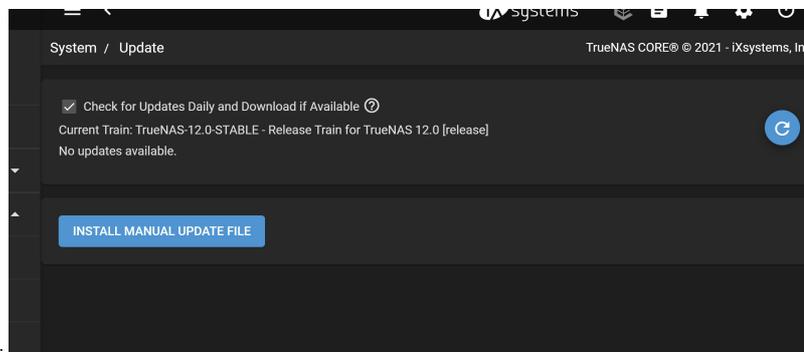
Once the installation completes, reboot the system and remove the USB with the SCALE .iso file.

When TrueNAS SCALE boots, you might need to [use the Shell to configure networking interfaces](#) to enable GUI accessibility.

Manual Update File

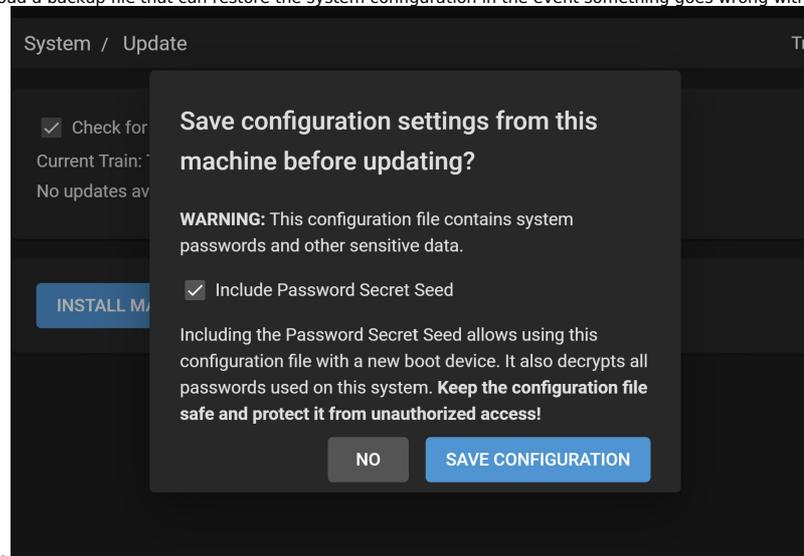
Start by downloading the [SCALE manual update file](#). Confirm that the TrueNAS system is on the latest public, 12.0-U8 or better, release.

Click **CHECK FOR UPDATES** in the **System Information** card on the **Dashboard** or go to **System > Update**.



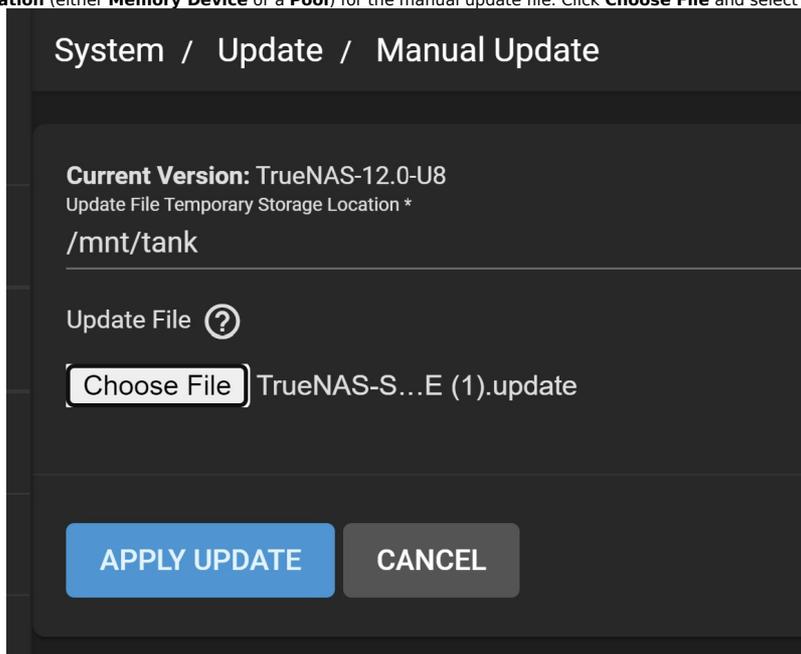
Click **INSTALL MANUAL UPDATE FILE**.

Click **SAVE CONFIGURATION** to download a backup file that can restore the system configuration in the event something goes wrong with the migration.



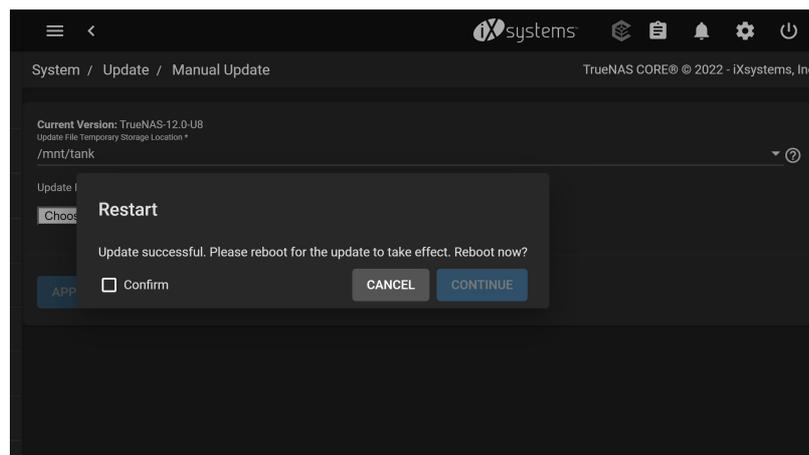
This is recommended but it not required.

Select a **Temporary Storage Location** (either **Memory Device** or a **Pool**) for the manual update file. Click **Choose File** and select the TrueNAS-



SCALE.update file you downloaded.

Then click **APPLY UPDATE**.



After the update completes, reboot the system.

Parallel SCALE CLI Commands

The following CLI commands are available after migrating from CORE to SCALE. The CORE equivalent CLI command is provided for reference. These commands are intended for diagnostic use. Making configuration changes using the SCALE OS CLI is not recommended.

CORE CLI Command	SCALE CLI Command	Description
camcontrol devlist	lshw -class disk -short sfdisk -l	Use <code>lshw -class disk -short sfdisk -l</code> to get detailed information on hardware (disk) configuration that includes memory, mainboard and cache configuration, firmware version, CPU version and speed.
geom disk list	lsblk, hdparm	Use <code>lsblk</code> to lists block devices or <code>hdparm</code> to get or set SATA/IDE device parameters.
glabel status	blkid	Use <code>blkid</code> to locate or print block device attributes.
gstat gstat -pods	iostat iostat -dtx	Use <code>iostat -dtx</code> to display the device utilization report with the time for each report displayed and includes extended statistics.
ifconfig ifconfig -l	ip addr ifconfig -s lshw -class network -short ethtool devname	Use <code>ip addr</code> to show or manipulate routing, devices, or policy routing and tunnels. Use <code>ifconfig -s</code> to configure a network interface. Use <code>lshw -class network -short</code> to display a network device tree showing hardware paths. Use <code>ethtool *devnam*</code> to query or control network driver and hardware settings.
netstat -i	ifstat -i	Use <code>ifstat -i</code> to get interface statistics on a list of interfaces to monitor.
nvmecontrol devlist	nvme list	Use <code>nvme list</code> to identify the list of NVMe devices on your system.
pmcstat	profile-bpfcc	Use <code>profile-bpfcc</code> to get a CPU usage profile obtained by sampling stack traces.
systat -ifstat	iftop netstat	Use <code>iftop</code> to display interface bandwidth usage by host and <code>netstat</code> to print network connections, routing tables, interface statistics, masquerade connections, and multicast memberships.
top -SHlZP	top -Hi	Use <code>top -Hi</code> to display Linux tasks for all individual threads and starts with the last remembered <i>i</i> state reversed.
vmstat -P	sar -P ALL	Use <code>sar -P ALL</code> to get reports with statistics for each individual processor and global statistics among all processors.

3.3 - Post-install Configuration

- [Configuring Network Interfaces](#)
- [Changing the Root Password](#)
- [Resetting the System Configuration](#)

The Console Setup menu displays at the end of the boot process. If the TrueNAS system has a keyboard and monitor, you can use this menu to administer the system.

By default, TrueNAS doesn't display the Console Setup menu when you connect via SSH or the web shell. The `root` user or another user with root permissions can start the Console Setup menu by entering `/etc/netcli`.

```
Console setup
-----
The web user interface is at:
http://10.0.2.15
https://10.0.2.15

1) Configure network interfaces
2) Reset root password
3) Reset configuration to defaults
4) Open TrueNAS CLI Shell
5) Open Linux Shell
6) Reboot
7) Shutdown

Enter an option from 1-7:
```

The menu provides these options:

1. **Configure network interfaces** provides options to set up network interfaces.
2. **Reset root password** resets the `root` user password.
3. **Reset configuration to defaults** resets the system to default settings.
4. **Open TrueNAS CLI Shell** starts a shell for running TrueNAS commands. Type `exit` to leave the shell.
5. **Open Linux Shell** starts a shell for running Linux commands. Type `exit` to leave the shell.
6. **Reboot** reboots the system.
7. **Shut Down** shuts down the system.

Console Setup menu options may change with software updates, service agreements, etc.

During boot, TrueNAS attempts to connect to a DHCP server from all live interfaces. If it receives an IP address, the Console Setup menu displays it so you can access the Web UI. In the picture above, the TrueNAS web UI is at `10.0.2.15`.

You may be able to access the web UI using `hostname.domain` (default is `truenas.local`) if your system:

- Doesn't have a monitor.
- Is on a network that supports Multicast DNS (mDNS).

Configuring Network Interfaces

Simple Network Interface Configuration

Simple Network Configuration

Enter 4 to open the TrueNAS CLI Shell.

```
Type "ls" (followed by Enter) to list available configuration options
[truenas]: network interface
Type "ls" (followed by Enter) to list available network interface configuration options
Type "man" (followed by Action) to get help on how to use the specific Action.

i.e.
man create
[truenas] network interface: query
-----
| name | type | state,aliases | aliases | ipv4_dhcp | ipv6_auto | description | mtu |
-----
| eno1 | PHYSICAL | fe80::3e0c:efff:fe5e:e772/64 | <empty list> | true | false | <null> | <null> |
| eno2 | PHYSICAL | <empty list> | <empty list> | true | false | <null> | <null> |
-----
[truenas] network interface: _
```

Select an Interface

1. Enter `network interface`, then enter `query` to display available physical system interfaces.
2. Once you know which interface you want to update, enter `update interface aliases=["ipaddress"] ipv4_dhcp=false`.
Example: `update eno1 aliases=["10.0.2.15"] ipv4_dhcp=false`
3. Enter `commit` to apply the pending changes.
4. Enter `checkin`, then enter `query` to show the updated interfaces.

```
[truenas] network interface: update eno1 aliases=["10.0.2.15/35"] ipv4_dhcp=false
You have pending network interface changes. Please run 'network interface commit'
to apply them.
[truenas] network interface: commit
<null>
Network interface changes have been applied. Please run 'network interface checkin'
if the network is still operational or they will be rolled back in 49 seconds.
[truenas] network interface: checkin
<null>
[truenas] network interface: query
-----
| name | type | state,aliases | aliases | ipv4_dhcp | ipv6_auto | description | mtu | disable_offload_capabilities |
-----
| eno1 | PHYSICAL | 10.0.2.15/35 | 10.0.2.15/35 | false | false | <null> | <null> | false |
| eno2 | PHYSICAL | <empty list> | <empty list> | false | false | <null> | <null> | undefined |
-----
[truenas] network interface:
```

Configure the Default Gateway

1. Enter `..` to go back to the `network>` prompt, then enter `configuration`.
2. Enter `update ipv4gateway="ipaddress"`. After you execute the command, the Console Setup menu displays the new web UI address.
Example: `update ipv4gateway="10.0.2.15"`

```
[truenas] network interface: ..
Type "ls" (followed by Enter) to list available network configuration options
[truenas] network configuration
Type "ls" (followed by Enter) to list available network configuration configuration options
Type "man" (followed by Action) to get help on how to use the specific Action.

i.e.
man activity_choices
[truenas] network configuration: update ipv4gateway="10.0.2.15/35"
[truenas] network configuration:
The web user interface is at:
http://10.0.2.15/35
https://10.0.2.15/35
```

3. Enter `exit` to go back to the Console Setup menu.

Configuring LAGG and VLAN

Configure LAGG

1. Enter 4 to open the TrueNAS CLI Shell.
2. Enter network interface, then enter query to display available physical system interfaces.
3. Once you know the interface names, enter create type=LINK_AGGREGATION lag_ports=["interface1","interface2"] lag_protocol=LACP
Example: network interface create type=LINK_AGGREGATION lag_ports=["eno1","eno2"] lag_protocol=LACP

```
[truenas] network interface: create type=LINK_AGGREGATION lag_ports=["eno1","eno2"] lag_protocol=LACP
You have pending network interface changes. Please run 'network interface commit' to apply them.
[truenas] network interface: commit
```

Configure VLAN

1. Enter create type=VLAN vlan_parent_interface=bond# vlan_tag=### aliases=[{"address": "ipaddress", "netmask": "bitlength"}]
Example: create type=VLAN vlan_parent_interface=bond0 vlan_tag=1022 aliases=[{"address": "10.0.2.15", "netmask": "32"}]

```
[truenas] network interface: create type=VLAN vlan_parent_interface=bond0 vlan_tag=1022 aliases=[{"address": "10.0.2.15", "netmask": "32"}]
You have pending network interface changes. Please run 'network interface commit' to apply them.
[truenas] network interface: commit
```

2. Enter commit to apply the pending changes.
3. Enter exit to return to the Console Setup menu.
4. Enter 5 to open the Linux Shell, then enter ip addr show to ensure you set the correct IP address.
5. Enter exit to go back to the Console Setup menu.

Configure Gateway

1. Enter 4 to open the TrueNAS CLI Shell.
2. Enter network configuration update ipv4gateway="ipaddress"
Example: network configuration update ipv4gateway="10.0.2.15"

```
[truenas]: network configuration update ipv4gateway="10.0.2.15"
[truenas]:
The web user interface is at:
http://10.0.2.15
https://10.0.2.15
```

3. Enter exit to go back to the Console Setup menu.
4. Enter 5 to open the Linux Shell.
5. Enter ping ipaddress to ping the gateway.
Example: ping 10.0.2.15
6. When you are ready to stop pinging, type Ctrl+C to view the statistics.

```
root@truenas[~]# ping 10.0.2.15
PING 10.0.2.15 (10.0.2.15) 56(84) bytes of data:
64 bytes from 10.0.2.15: icmp_seq=1 ttl=64 time=0.019 ms
64 bytes from 10.0.2.15: icmp_seq=2 ttl=64 time=0.029 ms
64 bytes from 10.0.2.15: icmp_seq=3 ttl=64 time=0.034 ms
^C
--- 10.0.2.15 ping statistics ---
 3 packets transmitted, 3 received, 0% packet loss, time 2030ms
rtt min/avg/max/mdev = 0.019/0.027/0.034/0.006 ms
root@truenas[~]#
```

7. Enter exit to go back to the Console Setup menu.

Configuring Static Route

1. Enter 4 to open the TrueNAS CLI Shell, then enter network interface
2. Enter update interface1 aliases="ipaddress"
Example: update eno1 aliases="10.0.2.15"

```
[truenas] network interface: update eno1 aliases="10.0.2.15"
You have pending network interface changes. Please run 'network interface commit' to apply them.
[truenas] network interface: commit
[null]
Network interface changes have been applied. Please run 'network interface checkin' if the network is still operational; or they will be rolled back in 48 seconds.
[truenas] network interface: checkin
[null]
[truenas] network interface:
```

3. Enter commit, then checkin to apply the changes.
4. Enter exit to go back to the Console Setup menu.

Changing the Root Password

Enter 2 in the Console Setup menu, then enter and re-enter the new password you want to use.

Changing the root password disables 2FA (Two-Factor Authentication).

Resetting the System Configuration

Enter 3 in the Console Setup menu, then enter y to reset the system configuration. The system will reboot and revert to default settings.

Caution! Resetting the configuration deletes all settings and reverts TrueNAS to default settings. Before resetting the system, back up all data and encryption keys/passphrases! After the system resets and reboots, you can go to **Storage** and click *Import* to re-import pools.

3.4 - First Time Login

Now that you have installed and configured TrueNAS SCALE, you can log in to the web interface and begin managing data!

Can I configure TrueNAS SCALE using a CLI?

After installing TrueNAS, you can configure and use the system through the web interface.

Important: Use only the web interface to make configuration changes to the system.

By default, using the command-line interface (CLI) to modify the system **does not modify the settings database**. The system reverts to the original database settings when it restarts and wipes any user-made command line changes. TrueNAS automatically creates several ways to access the web interface, but you might need to adjust the default settings for your network environment.

Web Interface Access

By default, fresh installs of TrueNAS SCALE provide a default address for logging in to the web interface. To view the web interface IP address or reconfigure web interface access, connect a monitor and keyboard to your TrueNAS system or connect with IPMI for out-of-band system management.

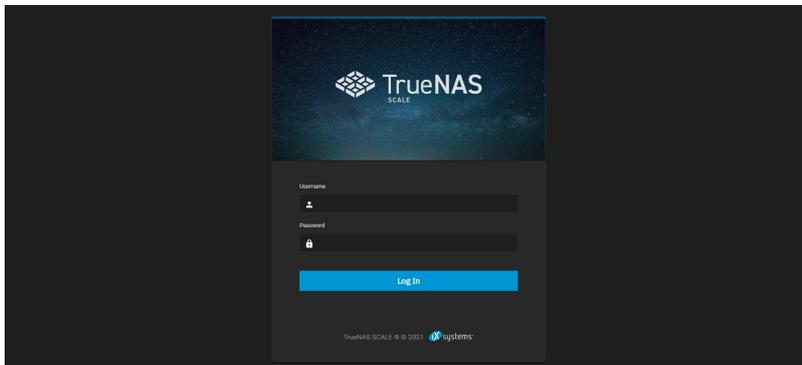
When powering on a TrueNAS system, the system attempts to connect to a DHCP server from all live interfaces to access the web UI. On networks that support Multicast Domain Name Services (mDNS), the system can use a hostname and domain to access the TrueNAS web interface. By default, TrueNAS uses the hostname and domain *truenas.local*. To change the hostname and domain in the web interface, go to **Network** and click *Settings* in the *Global Configuration* window.

To access the web interface using an IP address, use the one that the Console Setup Menu generated after installing SCALE, or use the one you configured in the [Post-install Configuration article](#) if you upgraded from CORE.

Create a strong login password! You can reset the root password in the TrueNAS console setup menu or web interface by going to **Credentials > Local Users** and editing the *root* user.

Logging In

On a computer with access to the same network as the TrueNAS system, enter the hostname and domain or IP address in a web browser to connect to the web interface.



Enter the *root* username and account password that you created during installation.

Troubleshooting

If the user interface is not accessible by IP address from a browser, check these things:

- If the browser configuration has proxy settings enabled, disable them and try connecting again.
- If the page does not load, ensure a ping reaches the TrueNAS system IP address. If the IP address is in a private range, you must access it from within that private network.

If the web interface displays but seems unresponsive or incomplete:

- Make sure the browser allows cookies, Javascript, and custom fonts from the TrueNAS system.
- Try a different browser. We recommend Firefox.

If the UI becomes unresponsive after an upgrade or other system operation, clear the site data and refresh the browser (Shift+F5).

Dashboard

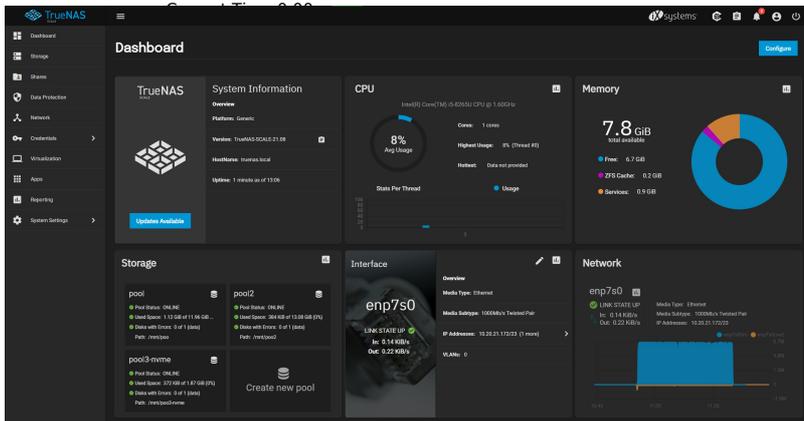


Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaletour.mp4>

Play Video

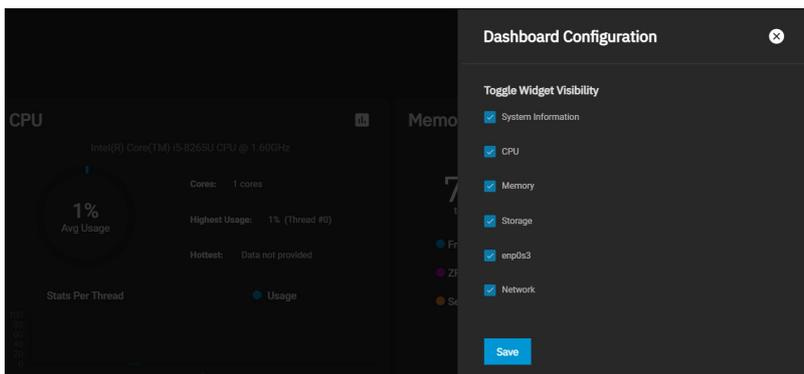
After logging in, you will see the system **Dashboard**. The dashboard displays basic information about the installed version, systems component usage, and network traffic. For users with compatible TrueNAS Hardware, clicking the system image will take you to the **System Settings > Enclosure** page.

Play



The **Dashboard** provides access to all TrueNAS management options. The top row has links to outside resources and buttons to control the system. The left-hand column lets users navigate to the various TrueNAS Configuration screens.

Users can select which widgets appear on the dashboard by clicking **Configure**.



Top Bar Menu

100%

Buttons in the top bar **TrueCommand** style the iXsystems site, display the status of TrueCommand, and show system processes and configuration menus.



iXsystems

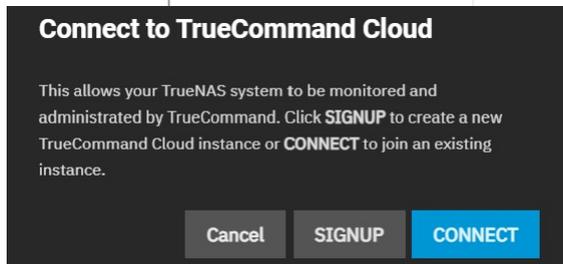
Close Modal Dialog

The iXsystems button opens the [iXsystems home page](#) where users can find information about storage and server systems.

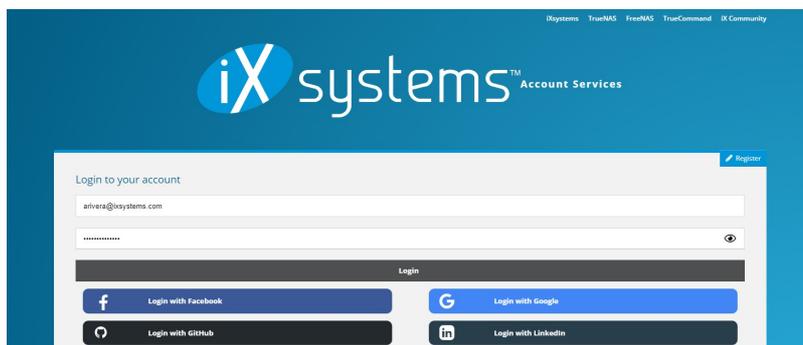
Users can also use the iXsystems home page to access their customer portal and community section for support.

Status of TrueCommand

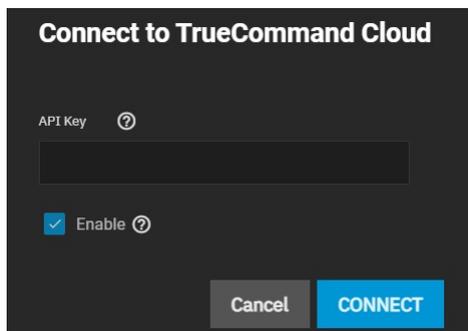
The *Status of TrueCommand* button lets users sign up with and connect to [TrueCommand Cloud](#).



Clicking *SIGNUP* will open the TrueCommand signup page in a new tab.

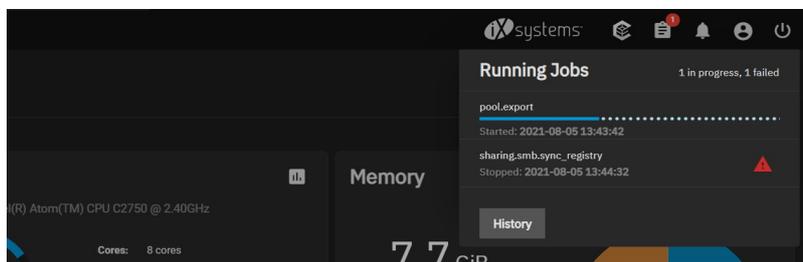


Once users have signed up, they can click the *CONNECT* button and enter their API key to connect SCALE to TrueCommand Cloud.



Task Manager

The Task Manager displays all running and failed jobs/processes.



Users can click the *History* button to open the **Jobs** screen. **Jobs** lists all *Successful*, *Active*, and *Failed* jobs. Users can also click *View Logs* next to a failed process to view its log information and error message.

Name	Date	ID	Started	Finished	Logs/Errors	Logs/Errors
sharing_smb.sync_registry	FAILED	135	2021-08-05 10:44:31	2021-08-05 10:44:32	View Logs	Log Path
gluster_eventd.delete	SUCCESS	134	2021-08-05 10:43:48	2021-08-05 10:43:48	None	Log Excerpt
gluster_eventd.init	SUCCESS	133	2021-08-05 10:43:48	2021-08-05 10:43:48	None	No logs are available
alert_process.alerts	SUCCESS	180	2021-08-05 14:06:42	2021-08-05 14:06:42	None	Error
alert_mail.alerts	SUCCESS	179	2021-08-05 11:05:42	2021-08-05 11:05:42	None	[ERR] [2] test conf taskname [found] failed with error: directory_create_or_atomic_readdir failed on directory /etc/dyptools/samba4/private/msg.sock: No such file or directory. Unable to initialize messaging context!
catalog.items	SUCCESS	93	2021-08-05 09:40:24	2021-08-05 09:40:27	None	

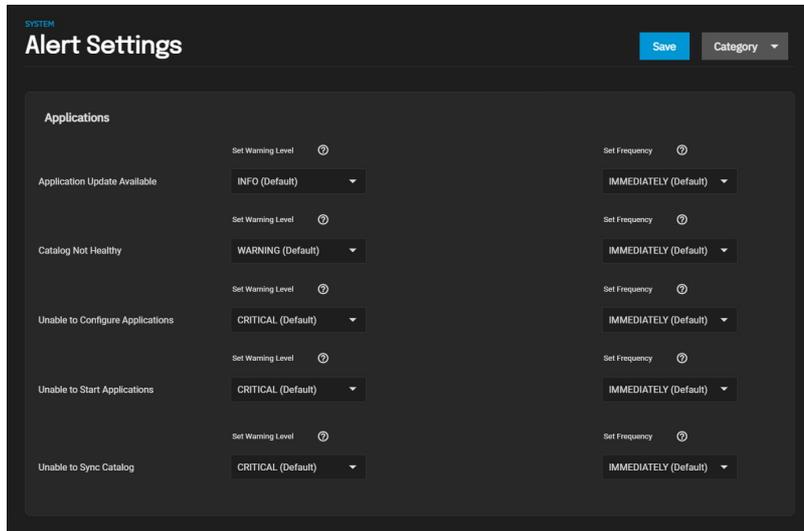
Alerts

The *Alerts* button displays the *Alerts* menu, which shows all current alerts. Users may dismiss alerts individually, or all at once.

The *Alerts* menu also lets users configure *Alert Settings*, *Alert Services*, and *Email*.

Alert Settings

The *Alert Settings* screen has options for setting the warning level and frequency for alerts specific to application actions.



The *Set Warning Level* drop-downs customize alert importance. Each warning level has an icon and color to express its urgency.

The *Set Frequency* drop-downs adjust how often the system sends alert notifications. Setting the *Frequency* to *NEVER* prevents that alert from being in the *Alerts* menu, but it will still pop up in the UI if triggered.

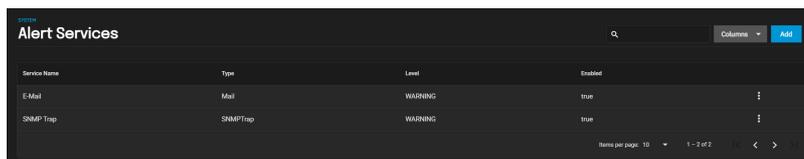
Alert Warning Levels

Each warning level has a different icon and color to express its urgency. To make the system email you when alerts with a specific warning level trigger, set up an email Alert Service with that warning level.

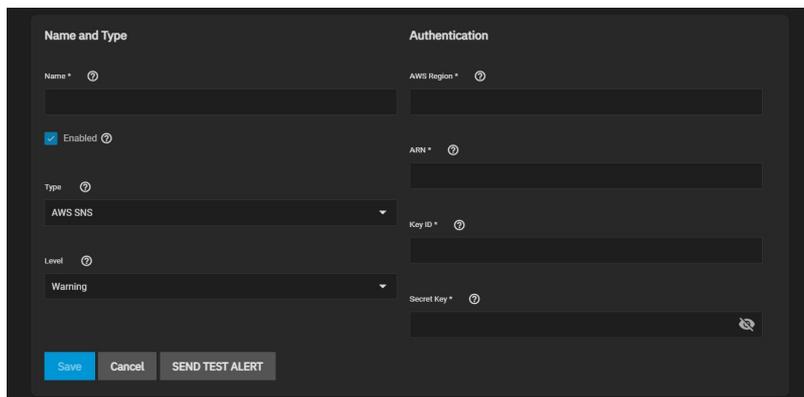
Level	Icon	Alert Notification?
1 INFO		No
2 NOTICE		Yes
3 WARNING		Yes
4 ERROR		Yes
5 CRITICAL		Yes
6 ALERT		Yes
7 EMERGENCY		Yes

Alert Services

The *Alert Services* screen has options to create and edit alert services. The *Alert Services* screen displays existing alert services in a list that users can filter by *Type*, *Level*, and *Enabled*.



To create a new alert service, click *Add* and fill out the form, then click *Save*.



Name and Type

Setting	Description
Name	Name of the new alert service.
Enabled	Unset to disable this service without deleting it.
Type	Choose an alert service to display options for that service.
Level	Select the level of severity.

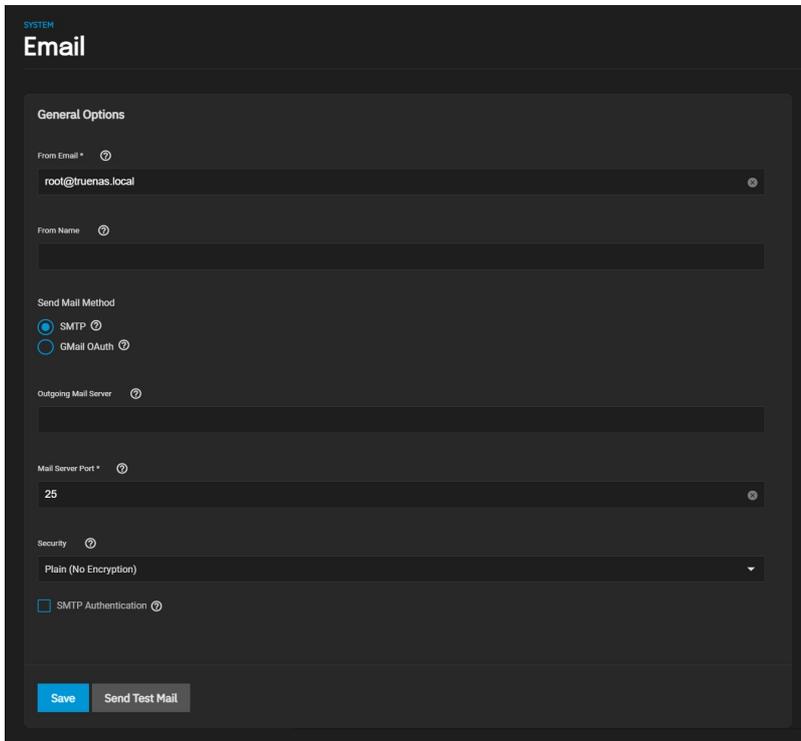
Authentication

Setting	Description
AWS Region	Enter the AWS account region .
ARN	Topic Amazon Resource Name (ARN) for publishing. Example: arn:aws:sns:us-west-2:111122223333:MyTopic.
Key ID	Access Key ID for the linked AWS account.
Secret Key	Secret Access Key for the linked AWS account.

The **SEND TEST ALERT** button generates a test alert to confirm the alert service will work correctly.

Email

The *Email* screen has options for users to set up a system email address.



Setting	Description
From Email	The user account Email address to use for the envelope From email address. The user account Email in Accounts > Users > Edit must be configured first.
From Name	The friendly name to show in front of the sending email address. Example: <i>Storage System 01</i> it@example.com
SMTP	Enable SMTP configuration.
GMail OAuth	Enable GMail OAuth authentication.
Outgoing Mail Server	Hostname or IP address of SMTP server to use for sending this email.
Mail Server Port	MTP port number. Typically 25,465 (secure SMTP), or 587 (submission).
Security	Email encryption type. Choices are Plain (No Encryption), SSL (Implicit TLS), or TLS (STARTTLS).
SMTP Authentication	Enable SMTP AUTH using PLAIN SASL. Requires a valid Username and Password.

The *Send Test Mail* button generates a test email to confirm the system email works correctly.

Settings

The *Settings* button has options for passwords, web interface preferences, API Keys, and TrueNAS information.

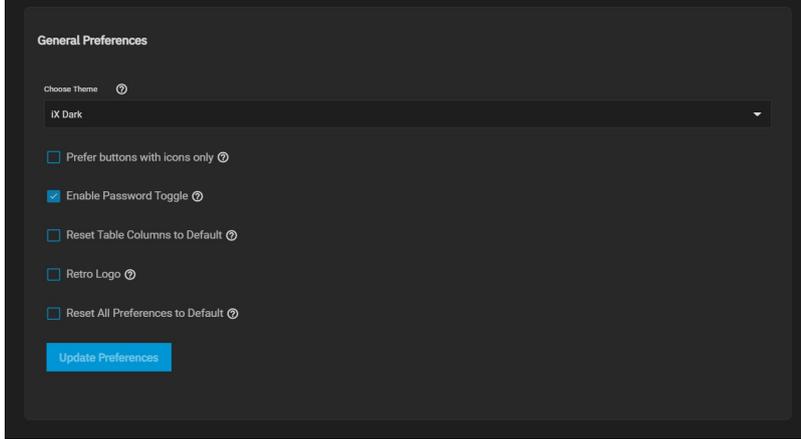
Change Password

Clicking **Change Password** allows users to change the currently logged-in administrator password.

Preferences

Clicking **Preferences** lets users select general preferences for the system.

Web Interface Preferences



Setting	Description
Choose Theme	Choose a preferred theme.
Prefer buttons with icons only	Preserve screen space with icons and tooltips instead of text labels.
Enable Password Toggle	When set, an eye icon appears next to password fields. Clicking the icon reveals the password.
Reset Table Columns to Default	Reset all tables to display default columns.
Retro Logo	Revert branding back to FreeNAS.
Reset All Preferences to Default	Reset all user preferences to their default values (custom themes are preserved).

API Keys

The *API Keys* section lets users add API Keys that identify outside resources and applications without a principal.

Users may also click *DOCS* to access their system's API documentation.

Guide and About

Clicking the *Guide* button opens the TrueNAS Documentation Hub in a new tab.

Clicking the *About* button brings up links to the TrueNAS Documentation Hub, the TrueNAS Community Forums, the FreeNAS Open Source Storage Appliance GitHub repository, and the iXsystems homepage.

Power

The *Power* button lets the user log out of, restart, or shut down the system.

Storing Data

Now that you can access the TrueNAS web interface and see all the management options, it's time to begin [storing data](#)!

4 - Storage

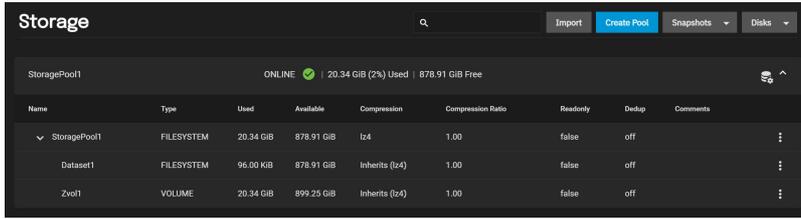
- [Storage Overview](#)

The SCALE Storage section has controls for pool, snapshot, and disk management.

The storage section also has options for datasets, Zvols, and permissions.

SCALE supports clustering storage across multiple systems. See [TrueCommand Clustering](#) for more details.

Storage Overview



The top row of the SCALE storage screen lets users search for existing pools, datasets, and zvols.

The *Import* button lets users reconnect pools exported/disconnected from the current system or created on another system. The import button also reconnects pools after users reinstall or upgrade the TrueNAS system.

The *Create Pool* button creates ZFS data storage “pools” with physical disks to efficiently store and protect data.

The *Snapshots* drop-down creates snapshots, which provide read-only point-in-time copies of a file system, volume, or a running virtual machine.

The *Disks* drop-down lets users manage, wipe, and import storage disks that TrueNAS will use for ZFS data storage.

The Storage screen displays the pools, datasets, and zvols users have created on the system. Users may perform actions to root pools or specific datasets using the *Pool Actions* and *Dataset Actions* menus.

Ready to get started? Choose a topic or article from the left-side Navigation pane. Click the < symbol to expand the menu to show the topics under this section.

4.1 - Pools

- [Review Storage Needs](#)
- [Creating or Importing a Pool](#)
- [Pool Operations](#)
 - [Pool Actions](#)
- [Encryption Actions](#)

TrueNAS uses ZFS data storage “pools” to efficiently store and protect data.

What's a pool?

Storage pools are attached drives organized into virtual devices (*vdevs*). ZFS and TrueNAS periodically review and *heal* when discovering a bad block in a pool. Drives are arranged inside *vdevs* to provide varying amounts of redundancy and performance. Combined, ZFS and *vdevs* combined create high-performance pools, pools that maximize data lifetime, and all situations in between.

Review Storage Needs

We strongly recommend users review the available system resources and plan the storage use case before creating a storage pool.

- Allocating more drives to a pool increases redundancy when storing critical information.
- Maximizing total available storage at the expense of redundancy or performance entails allocating large-volume disks and configuring a pool for minimal redundancy.
- Maximizing pool performance entails installing and allocating high-speed SSD drives to a pool.

Determining your specific storage requirements is a critical step before creating a pool.

Creating or Importing a Pool

Creating a Pool

Tutorial



Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaleangelfishpoolcreate.mp4>

Play Video

Play

Mute

Current Time 0:00

/

Duration 1:15

Loaded: 100.00%

Stream Type LIVE

Seek to live, currently behind liveLIVE

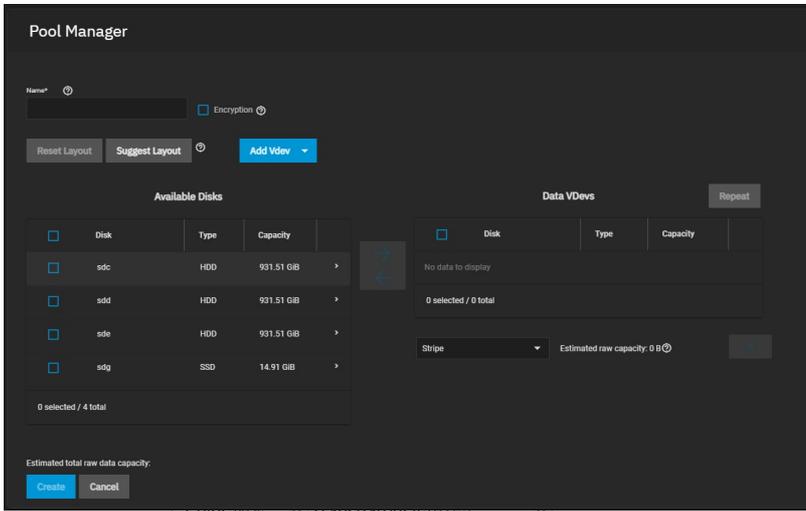
Remaining Time -1:15

1x

Playback Rate

Chapters

To create a new pool, go to **Storage** and click **Create Pool**.



First, enter a pool name.

Encryption? TrueNAS offers several encryption algorithms to maximize security. However, encryption also complicates data retrieval and risks permanent data loss! Refer to the [Encryption article](#) for more details and decide if encryption is necessary for your use case before setting any **Encryption** option.

Next, configure the virtual devices (vdevs) that make up the pool.

Suggested Layout

Clicking **Suggest Layout** allows TrueNAS to review all available disks and populate the primary **Data VDevs** with identically sized drives in a configuration balanced between storage capacity and data redundancy. Click **Reset Layout** to clear the suggestion.

To manually configure the pool, add vdevs according to your use case. Check the **Disk** checkboxes and click the → to move the disks into a vdev.

Vdev Types

Pools offer several vdev types. Vdevs store data or enable unique features for the pool.

To add a vdev type during pool creation, click **Add Vdev()** and select the type. Select disks from **Available Disks** and use the → (right arrow) next to the new VDev to add it to that section.

Data Type

Data is the standard vdev for primary storage operations. Each storage pool requires at least one Data vdev. **Data** vdev configuration typically affects how users can configure other types of vdevs.

Duplicating a Data vdev

Users can duplicate a **Data VDev** by clicking **Repeat**. When the system has more available equal-sized disks, the **Repeat** button creates another vdev with an identical configuration called a **Mirror**.

When even more same-size disks are available, users can create multiple copies of the original vdev.

We do not recommend having multiple data vdevs with different numbers of disks in each vdev. Doing so complicates and limits pool capabilities.

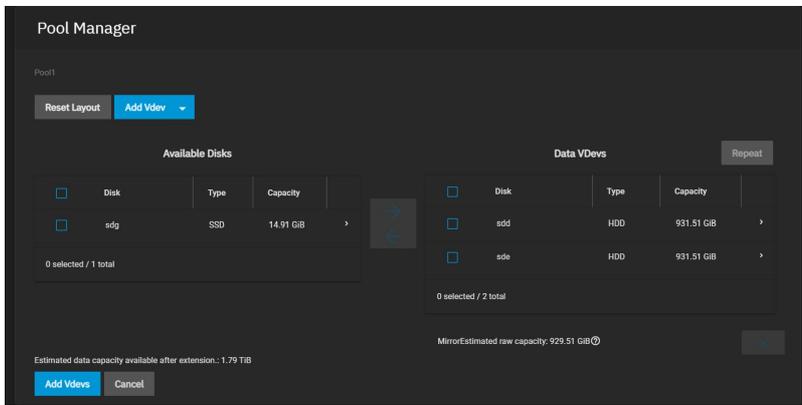
Additional Types

- **Cache:** [ZFS L2ARC](#) read-cache used with fast devices to accelerate read operations. Users can add or remove Cache VDevs after creating the pool.
- **Log:** [ZFS LOG](#) device that improves synchronous write speeds. Users can add or remove Log VDevs after creating the pool.
- **Hot Spare:** Drives reserved for inserting into **Data** vdevs when an active drive fails. The system uses hot spares as temporary replacements for failed drives to prevent larger pool and data loss scenarios. When a user replaces a failed drive with a new one, the hot spare reverts to an inactive state and becomes available again as a hot spare. If a user detaches the failed drive from the pool without adding a new one, the system promotes the temporary hot spare to a full Data vdev member.
- **Metadata:** Special allocation class used to create Fusion Pools for increased metadata and small block I/O performance.
- **Dedup:** Stores [ZFS de-duplication](#). Requires allocating X GiB for every X TiB of general storage. Example: 1 GiB of *Dedup* vdev capacity for every 1 TiB of *Data* vdev availability.

Vdev Layouts

Disks added to a vdev arrange in different layouts, according to the specific pool use case.

The **Pool Manager** suggests a vdev layout from the number of disks added to the vdev. For example, if you add two disks, TrueNAS automatically configures the vdev as a Mirror. The total available storage is the size of one added disk while the other disk provides redundancy.



To change the vdev layout, open the **Data VDevs** list and select the desired layout.

Can I create vdevs with different layouts in one pool?

TrueNAS SCALE does not support adding multiple vdevs with different layouts to a pool. Create a new pool when a different vdev layout is required. For example, *pool1* has a data vdev in a *mirror* layout, so create *pool2* for any *raid-z* vdevs.

- **Stripe**: Each disk stores data. A Stripe requires at least one disk and has no data redundancy.
- **Mirror**: Data is identical in each disk. A Mirror requires at least two disks, provides the most redundancy, and has the least capacity.
- **RAIDZ1**: Uses one disk for parity while all other disks store data. RAIDZ1 requires at least three disks.
- **RAIDZ2**: Uses two disks for parity while all other disks store data. RAIDZ2 requires at least four disks.
- **RAIDZ3**: Uses three disks for parity while all other disks store data. RAIDZ3 requires at least five disks.

Never use **Stripe** to store critical data! A single disk failure results in losing all data in the vdev.

Importing a Pool

The import procedure only applies to disks with a ZFS storage pool. To import disks with different file systems, see the [SCALE Disks](#) article.

ZFS pool importing works for pools that were exported or disconnected from the current system, created on another system, and pools to reconnect after reinstalling or upgrading the TrueNAS system.

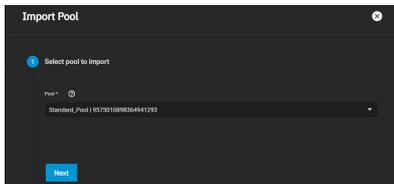
Do I need to do anything different with disks installed on a different system?

When physically installing ZFS pool disks from another system, use the `zpool export poolname` command in the command line or a web interface equivalent to export the pool on that system. Shut that system down and move the drives to the TrueNAS system. Shutting down the original system prevents an **in use by another machine** error during the TrueNAS import.

To import a pool, go to **Storage** and click **Import**.

TrueNAS detects any pools that are present but unconnected.

Select a pool from the **Pool** drop-down and click **Next**.



Review the Pool Import Summary and click **Import**.

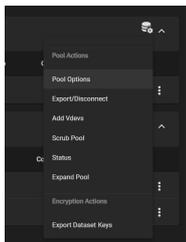


Can I import GELI-encrypted pools?

Since GELI encryption is specific to FreeBSD, TrueNAS SCALE cannot import GELI-encrypted pools. See the **Migrating GELI-encrypted Pools to SCALE** section in the [Installing SCALE](#) article.

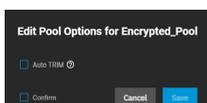
Pool Operations

Use the **Pool Operations** button to manage a pool.



Pool Actions

Pool Options

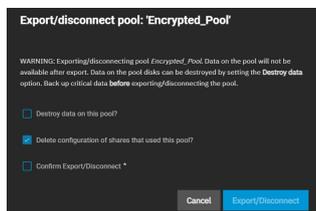


Pool Options allows users to set/unset **Auto TRIM**.

With **Auto TRIM** selected and active, TrueNAS periodically checks the pool disks for storage blocks it can reclaim. Auto TRIM can impact pool performance, so it's disabled by default.

For more details about *TRIM* in ZFS, see the `autotrim` property description in [zpool.8](#).

Export/Disconnect



The **Export/Disconnect** option disconnects the pool to transfer drives to a new system for importing or completely deletes the pool and any data stored on it.

A dialog box displays any system services affected by exporting the pool.

Users can erase all data on the pool by checking the **Destroy data on this pool?** box.

Clicking the **Delete configuration of shares that used this pool?** box deletes shares connected to the pool.

Add Vdevs

The **Add Vdevs** button opens the **Pool Manager** so users can add Vdevs to the pool.

Users cannot change the original encryption or data Vdev configuration.

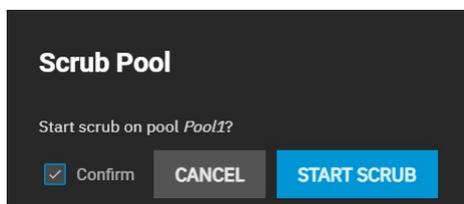
TrueNAS selects **Data VDev** by default. To add different Vdev types to the pool, select one from the **Add Vdev** drop-down.

When adding disks to increase the capacity of a pool, ZFS supports adding Vdevs (virtual devices) to an existing ZFS pool. After creating a Vdev, you cannot add more drives to that Vdev. However, you can stripe a new Vdev with another of the same type to increase the overall pool size. To extend a pool, you must add a Vdev that is the same type as existing Vdevs.

Vdevs extending examples:

- To make a striped mirror, add the same number of drives to extend a ZFS mirror. For instance, if ten new drives are available, you could initially create a mirror of two drives, then extend the mirror by adding another mirror of two drives, repeating three more times until you've added all ten drives.
- To make a stripe of two RAIDZ1 Vdevs (similar to RAID 50 on a hardware controller), add another three drives to extend a three-drive RAIDZ1.
- To make a stripe of RAIDZ2 vdevs (similar to RAID 60 on a hardware controller), add another four drives to extend a four-drive RAIDZ2.
- Add a disk as a Hot Spare to the pool.

Scrub Pool



Scrub Pool initiates a pool data integrity check.

If TrueNAS detects any problems during the scrub, it either corrects them automatically or generates an [alert](#) in the web interface.

By default, TrueNAS automatically checks every pool is on a reoccurring [scrub schedule](#).

Status

Name	Read	Write	Checksum	Status
Encrypted_Pool	0	0	0	ONLINE
RAIDZ1	0	0	0	ONLINE

Status displays the state of the last scrub and disks in the pool.

The **Pool Status** screen has additional [disk management](#) options.

Expand Pool

Expand Pool increases the pool size to match all available disk space. A user typically expands a pool when virtual disks are resized apart from TrueNAS.

Upgrade Pool

The **Upgrade Pool** option only appears when TrueNAS can upgrade the pool to use new [ZFS feature flags](#).

Before upgrading an existing pool, be aware of these caveats:

- Upgrading a pool is one-way. You cannot regress to an earlier ZFS version or downgrade to an earlier software version that does not support those ZFS features.
- Upgrading a pool unlikely affects data, but we recommend backing up data for safety. Before performing any operation that can affect the data on a storage disk, always back up all data first and verify the backup's integrity.
- Upgrading a ZFS pool is optional. Do not upgrade the pool if you may revert to an earlier TrueNAS version or repurpose the disks in another OS that supports ZFS. Upgrading a pool is unnecessary unless the end-user specifically needs the newer ZFS Feature Flags. If you upgrade a pool to the latest feature flags, you cannot import it into another OS that does not yet support them.

Upgrading a pool only takes a few seconds and is non-disruptive. You do not need to stop any sharing services to upgrade a pool. However, we recommend upgrading when the pool is not seeing heavy use. Upgrading suspends I/O for a short period, but is nearly instantaneous on a quiet pool.

Encryption Actions

See the [SCALE Encryption](#) page for detailed encryption information.

4.1.1 - Datasets

- [Creating a Dataset](#)
 - [Dataset Options](#)
- [Managing Datasets](#)
 - [Quotas](#)

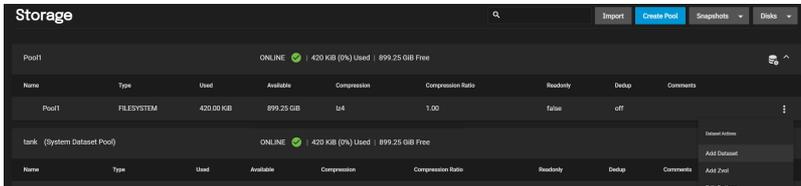
A TrueNAS dataset is a file system within a data storage pool. Datasets can contain files, directories (child datasets), and have individual permissions or flags. Datasets can also be [encrypted](#), either using the encryption created with the pool or with a separate encryption configuration.

We recommend organizing your pool with datasets before configuring [data sharing](#), as this allows for more fine-tuning of access permissions and using different sharing protocols.

Creating a Dataset

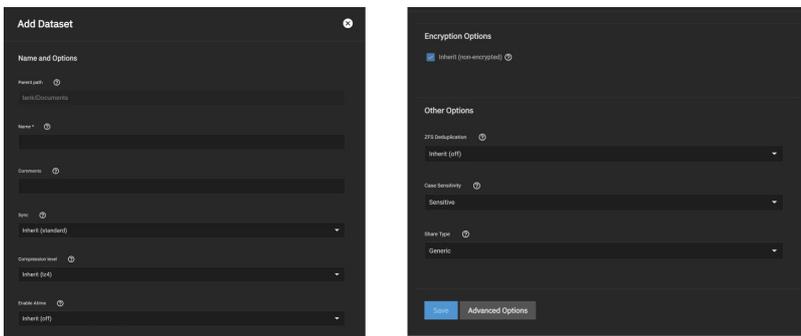
To create a dataset in the desired pool, go to **Storage**.

Select the pool top-level (root) dataset, and then click the  and select **Add Dataset**.



To create a dataset with the default options, enter a name and click **Submit**.

Dataset Options



You must configure the **Name and Options** fields to create the dataset. Datasets typically inherit settings from the root or parent dataset, so you only need to enter the dataset **Name** before clicking **Submit**.

Basic Options

Setting	Description
Name	Unique identifier for the dataset. Cannot change after creating and saving the dataset.
Comments	Notes about the dataset.
Sync	Standard uses the sync settings requested by the client software. Always waits for data writes to complete, and Disabled never waits for writes to complete.
Compression level	A drop-down list.
Enable Atime	A drop-down list.

By default, datasets inherit the encryption settings on the *Encryption Options* form from the root or parent dataset. To configure the dataset with different encryption settings, clear the **Inherit** checkbox and enter the new **Encryption Options** form values. For detailed descriptions of the encryption options, see the [Encryption article](#).

Other Options

The **Other Options** help tune the dataset for specific data sharing protocols:

Setting	Description
ZFS Deduplication	Transparently reuse a single copy of duplicated data to save space. Deduplication can improve storage capacity, but is RAM intensive. Compressing data is recommended before using deduplication. Deduplicating data is a one-way process. <i>Deduplicated data cannot be unduplicated!</i>
Case Sensitivity	Sensitive assumes filenames are case sensitive. Insensitive assumes filenames are not case sensitive. Mixed understands both types of filenames. You cannot change this after the saving the dataset.
Share Type	Define the type of data sharing the dataset uses to optimize the dataset for that sharing protocol. You cannot change this setting after the saving dataset.

Advanced Options

Clicking **Advanced Options** adds dataset quota management tools and a few additional fields to the **Other Options**:

Setting a quota defines the maximum allowed space for the dataset. You can also reserve a defined amount of pool space to prevent automatically generated data like system logs from consuming all of the dataset space. You can configure quotas for only the new dataset or include all child datasets.

Setting	Description
Quota for this dataset	Define the maximum allowed space for the dataset. 0 disables quotas.
Quota warning alert at, %	Generate a warning level alert when consumed space reaches the defined percentage. By default, the dataset inherits this value from the parent dataset. Clear the Inherit checkbox to change the value.
Quota critical alert at, %	Generate a critical level alert when consumed space reaches the defined percentage. By default, the dataset inherits this value from the parent dataset. Clear the Inherit checkbox to change the value.
Reserved space for this dataset	Reserve additional space for datasets that contain logs which could eventually take up all the available free space. 0 is unlimited.

TrueNAS adds more fields to the **Other Options**. By default, many of these options inherit their values from the parent dataset.

Setting	Description
Read-only	On prevents modifying the dataset. Off allows users accessing the dataset to modify its contents.
Exec	On allows executing processes from within this dataset. Off prevents executing processes from with the dataset. We recommend setting it to On .

Snapshot directory	Controls visibility of the <code>.zfs</code> directory on the dataset. Select either Visible or Invisible .
Copies	Duplicates ZFS user data stored on this dataset. Choose between 1 , 2 , or 3 redundant data copies. This can improve data protection and retention, but is not a substitute for storage pools with disk redundancy.
Record Size	Logical block size in the dataset. Matching the fixed size of data, as in a database, can result in better performance.
ACL Type	Inherit preserves ACL type from the parent dataset. Off to use neither NFSv4 or POSIX protocols. NFSv4 use to losslessly migrate Windows-style ACLs across Active Directory domains (or stand-alone servers) that use ACL models richer than POSIX. Since POSIX ACLs are a Linux-specific ZFS feature, administrators should use NFSv4 to maintain compatibility with TrueNAS Core, FreeBSD, or other non-Linux ZFS implementations. POSIX use when an organization data backup target does not support native NFSv4 ACLs. Since the Linux platform used POSIX for a long time, many backup products that access the server outside the SMB protocol cannot understand or preserve native NFSv4 ACLs. <i>All datasets within an SMB share path must have identical ACL types.</i> For a more in-depth explanation of ACLs and configurations in TrueNAS SCALE, see our ACL Primer .
ACL Mode	Determine how <code>chmod</code> behaves when adjusting file ACLs. See the <code>zfs(8) acemode</code> property. Passthrough only updates ACL entries that are related to the file or directory mode. Restricted does not allow <code>chmod</code> to make changes to files or directories with a non-trivial ACL. An ACL is trivial if it can be fully expressed as a file mode without losing any access rules. Set the ACL Mode to restricted to optimize a dataset for SMB sharing, but it can require further optimizations. For example, configuring an rsync task with this dataset could require adding <code>--no-perms</code> in the task Auxiliary Parameters field.
Metadata (Special) Small Block Size	Threshold block size for including small file blocks into the special allocation class (fusion pools) . Blocks smaller than or equal to this value are assigned to the special allocation class while greater blocks are assigned to the regular class. Valid values are zero or a power of two from 512B up to 1M. The default size 0 means no small file blocks are allocated in the special class. Before setting this property, you must add a special class vdev to the pool.

Managing Datasets

After creating a dataset, users can manage additional options by going to **Storage** and selecting the dataset and clicking **i** for that dataset.

- **Add Dataset** creates a new *child* dataset in this dataset. You can continue to layer datasets in this manner.
- **Add Zvol** creates a new [ZFS block device](#) as a child in this dataset.
- **Edit Options** allows users to adjust the dataset configuration. Users cannot change the **Name**, **Case Sensitivity**, or **Share Type** after saving the dataset.
- **Edit Permissions** opens the dataset access permissions editor. Depending on the ACL type users select during dataset creation, the permissions editor is either a simple permissions editor or the full ACL editor. For more information, see the [permissions](#) article.
- **User Quotas** lets users set data or object quotas for user accounts cached on or connected to the system.
- **Group Quotas** lets users set data or object quotas for user groups cached on or connected to the system.
- **Delete Dataset** removes the dataset, all stored data, and any snapshots from TrueNAS.

Deleting datasets can result in unrecoverable data loss! Be sure to move or obsolete any critical data off the dataset.

- **Create Snapshot** takes a single dataset [ZFS snapshot](#) to provide additional data protection and mobility. The system lists created snapshots on the **Snapshots** screen.

Quotas

TrueNAS allows setting data or object quotas for user accounts and groups cached on or connected to the system.

User

To view and edit user quotas, go to **Storage** and click **i** next to a dataset to open the **Dataset Actions** menu, then select **User Quotas**.

Name	Data Quota	DQ Used	DQ % Used	Object Quota	Objects Used	OQ % Used
root	0 bytes	512,000 bytes	0%	0	1	0%

The **User Quotas** page displays the names and quota data of any user accounts cached on or connected to the system.

To edit individual user quotas, go to the user row and click **i**, then click **Edit User**.

The **Edit User** window lets users edit the **User Data Quota** and **User Object Quota** values.

User Data Quota is the amount of disk space that selected users can use. **User Object Quota** is the number of objects selected users can own.

To edit user quotas in bulk, click **Actions** and select **Set Quotas (Bulk)**.

The **Set Quotas** window lets you edit user data and object quotas after selecting any cached or connected users.

Group

Go to **Storage** and click **i** next to a dataset to open the **Dataset Actions** menu, then select **Group Quotas**.

Name	Data Quota	DD Bytes Used	DD % Used	Object Quota	DD Object Used	DD % Used
root	0 bytes	512	0%	0	1	0%

The **Group Quotas** page displays the names and quota data of any groups cached on or connected to the system.

To edit individual group quotas, go to the group row and click **>**, then click **Edit**.

The **Edit Group** window lets users edit the **Group Data Quota** and **Group Object Quota** values.

To edit group quotas in bulk, click **Actions** and select **Set Quotas (Bulk)**.

TrueNAS presents the same options for single groups and lets users choose groups for the new quota rules.

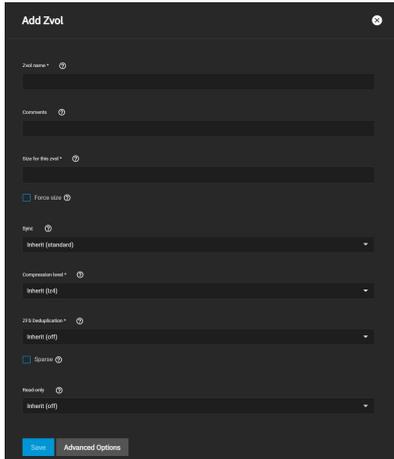
4.1.2 - Zvols

- [Zvol Creation Options](#)
- [Managing Zvols](#)

A ZFS Volume (zvol) is a [dataset](#) that represents a block device. TrueNAS requires a zvol when configuring [iSCSI Shares](#).

To create a zvol in a pool, go to **Storage** and click **+** on a pool or dataset, then select *Add Zvol*.

Zvol Creation Options



To create a zvol with default options, enter a name and size for the zvol and click *Save*.

Basic Options

Setting	Description
Zvol name	Enter a short name for the zvol. Using a zvol name longer than 63-characters can prevent accessing zvols as devices. For example, a zvol with a 70-character filename or path cannot be used as an iSCSI extent. This setting is mandatory.
Comments	Enter any notes about this zvol.
Size for this zvol	Specify size and value. Units like <code>t</code> , <code>TiB</code> , and <code>G</code> can be used. The size of the zvol can be increased later, but cannot be reduced. If the size is more than 80% of the available capacity, the creation will fail with an "out of space" error unless <i>Force size</i> is also enabled.
Force size	By default, the system will not create a zvol if that operation brings the pool to over 80% capacity. While NOT recommended , enabling this option will force the creation of the zvol.
Sync	Sets the data write synchronization. <i>Inherit</i> gets the sync settings from the parent dataset. <i>Standard</i> uses the sync settings requested by the client software. <i>Always</i> waits for data writes to complete. <i>Disabled</i> never waits for writes to complete.
Compression level	Compress data to save space. Refer to Compression for a description of the available algorithms.
ZFS Deduplication	Do not change this setting unless instructed to do so by your iXsystems support engineer.
Sparse	Used to provide thin provisioning. Use with caution as writes will fail when the pool is low on space.
Read-only	Set to prevent the zvol from being modified.
Inherit (Encryption Options)	Enabling causes the zvol to use the encryption properties of the root dataset.

Advanced Options

Setting	Description
Block size	The default is <i>Inherit</i> , other options include, <i>4KiB</i> , <i>8KiB</i> , <i>16KiB</i> , <i>32KiB</i> , <i>64KiB</i> , <i>128KiB</i>

TrueNAS automatically recommends a space-efficient *block size* for new zvols. This table shows the minimum recommended volume *block size* values. Use the *Block size* drop-down to change the value.

Configuration	Number of Drives	Optimal Block Size
Mirror	N/A	16k
Raidz-1	3	16k
Raidz-1	4/5	32k
Raidz-1	6/7/8/9	64k
Raidz-1	10+	128k
Raidz-2	4	16k
Raidz-2	5/6	32k
Raidz-2	7/8/9/10	64k
Raidz-2	11+	128k
Raidz-3	5	16k
Raidz-3	6/7	32k
Raidz-3	8/9/10/11	64k
Raidz-3	12+	128k

Depending on their workload, zvols can require additional tuning for optimal performance. See the OpenZFS handbook [workload tuning chapter](#) for more information.

Managing Zvols

To see zvol options, click **⋮** next to the desired zvol in **Storage**:

- *Delete Zvol* removes the zvol from TrueNAS. Deleting a zvol also deletes all of that zvol's snapshots.

Deleting zvols can result in unrecoverable data loss! Be sure that any critical data is moved off the zvol or is otherwise obsolete.

- *Edit Zvol* opens the zvol creation form for changing the previously saved settings. Users cannot change the name.
- *Create Snapshot* takes a single current point-in-time image of the zvol and saves it to *Snapshots*. TrueNAS will suggest a *Name* and offer the *Recursive* option.

If you clone a zvol from an existing [snapshot](#), TrueNAS will offer the *Promote Dataset* option. After promoting a clone, the original volume becomes a clone of

the promoted clone. Promoting a clone allows users to delete the volume that created the clone. Otherwise, you cannot delete a clone while the original volume exists.

When a zvol is the child of an [encrypted](#) dataset, TrueNAS offers additional *Encryption Actions*.

4.1.3 - Permissions

- [ACL Types in SCALE](#)
- [Unix Permissions Editor \(POSIX\)](#)
- [ACL Manager \(NFSv4\)](#)
 - [Permissions and Flags](#)

TrueNAS SCALE has a simple permissions manager and a full Access Control List (ACL) editor that defines dataset permissions. Permissions control the actions users can perform on dataset contents.

ACL Types in SCALE

TrueNAS SCALE offers two ACL types: POSIX (SCALE default) and NFSv4.

You can select which ACL types you want new datasets to use while creating them. To change an existing dataset's ACL type, click the [button](#) next to the intended dataset and select *Edit Options*. Next, click *Advanced Options* and scroll down to the *ACL Type* dropdown.

WARNING: Changing the ACL type affects how TrueNAS writes and reads on-disk ZFS ACL.

When the ACL type changes from POSIX to NFSv4, internal ZFS ACLs do not migrate by default, and access ACLs encoded in `posix1e` extended attributes convert to native ZFS ACLs.

When the ACL type changes from NFSv4 to POSIX, native ZFS ACLs do not convert to `posix1e` extended attributes, but ZFS will use the native ACL for access checks.

To prevent unexpected permissions behavior, you must manually set new dataset ACLs recursively after changing the ACL type.

Setting new ACLs recursively is destructive. We suggest creating a ZFS snapshot of the dataset before changing the ACL type or modifying permissions.

For a more in-depth explanation of ACLs and configurations in TrueNAS SCALE, see our [ACL Primer](#).

ACL Details from Shell

To view ACL information from the console, go to **System Settings > Shell** and enter:

```
getfacl /mnt/path/to/dataset
```

Unix Permissions Editor (POSIX)

The **Unix Permissions Editor** option allows basic adjustments to a dataset's ACL. Click [button](#) for the dataset to edit and choose *View Permissions*. Select [button](#) to access the **Unix Permissions Editor** page.

Field	Description
User	Select a user to control the dataset. Users created manually or imported from a directory service appear in the menu.
Apply User	Confirms changes to <i>User</i> . To prevent errors, TrueNAS only submits <i>User</i> changes when you set this box.
Group	Select the group to control the dataset. Groups created manually or imported from a directory service appear in the menu.
Apply Group	Confirms changes to <i>Group</i> . To prevent errors, TrueNAS only submits <i>Group</i> changes when you set this box.

Owner

The *Owner* section controls which TrueNAS *User* and *Group* has full control of this dataset.

Field	Description
User	Select a user to control the dataset. Users created manually or imported from a directory service appear in the menu.
Apply User	Confirms changes to <i>User</i> . To prevent errors, TrueNAS only submits <i>User</i> changes when you set this box.
Group	Select the group to control the dataset. Groups created manually or imported from a directory service appear in the menu.
Apply Group	Confirms changes to <i>Group</i> . To prevent errors, TrueNAS only submits <i>Group</i> changes when you set this box.

Access

The *Access* section lets users define the basic *Read*, *Write*, and *Execute* permissions for the *User*, *Group*, and *Other* accounts that might access this dataset.

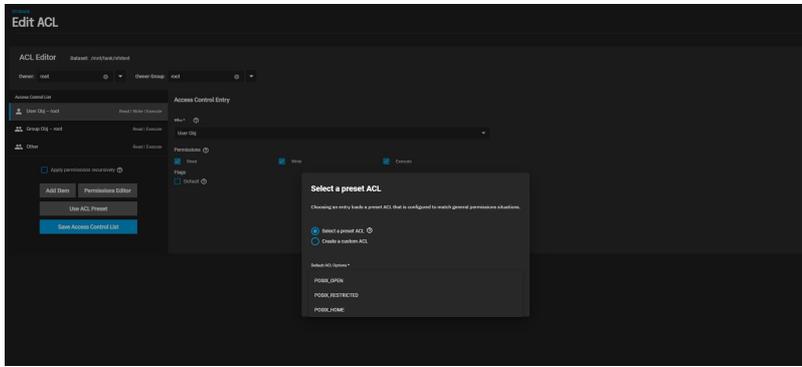
Advanced

The *Advanced* section lets users *Apply Permissions Recursively* to all directories, files, and child datasets within the current dataset. Users may also set *Traverse*, which applies permissions recursively to all child datasets in the current dataset.

To switch from the basic permissions editor to the advanced ACL editor, click *Set ACL*.

An Access Control List (ACL) is a set of account permissions associated with a dataset and applied to directories or files within that dataset. TrueNAS uses ACLs to manage user interactions with shared datasets and creates them when users add a dataset to a pool.

The TrueNAS has options to *Select a preset ACL* or *Create a custom ACL*. The available preset ACLs are *POSIX_OPEN*, *POSIX_RESTRICTED*, or *POSIX_HOME*.



When creating a custom ACL, use *Add Item* to apply additional permissions to the *Access Control List*.

POSIX ACL Editor

Field	Description
Path	Shows the full pathway to the file.
Owner	User who controls the dataset. This user always has permissions to read or write the ACL and read or write attributes. Users created manually or imported from a directory service appear in the drop-down menu.
Owner Group	The group which controls the dataset. This group has the same permissions as granted to the group@ Who. Groups created manually or imported from a directory service appear in the drop-down menu.

You can select any user accounts or groups imported from a directory service as the primary *User* or *Group*.

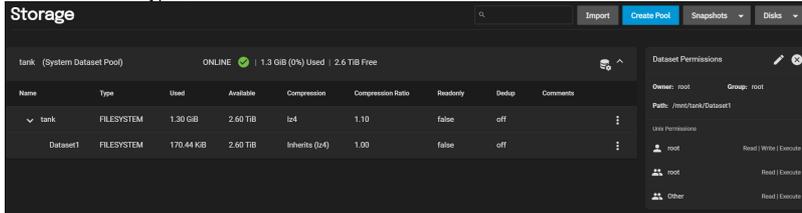
Access Control List

Define *Who* the Access Control Entry (ACE) applies to and configure permissions and inheritance flags for the ACE.

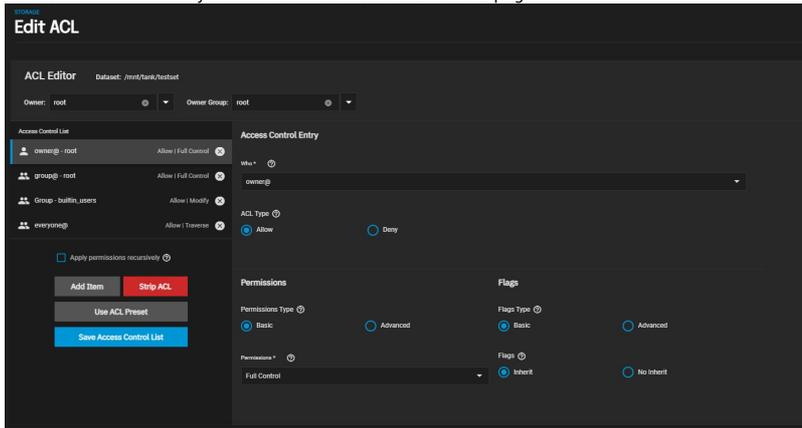
Field	Description
Who	Access Control Entry (ACE) user or group. Select a specific User or Group for this entry. Owner@ applies this entry to the user that owns the dataset. Group@ applies this entry to the group that owns the dataset. Everyone@ applies this entry to all users and groups. See nfs4_setfac(1) NFSv4 ACL ENTRIES .
User	User account to which this ACL entry applies.
Permissions	Select permissions to apply to the chosen Who. Choices change depending on the Permissions Type.
Flags	How this ACE applies to newly created directories and files within the dataset. Basic flags enable or disable ACE inheritance. Advanced flags allow further control of how the ACE applies to files and directories in the dataset.

ACL Manager (NFSv4)

For *NFSv4* share types, click [here](#) for the dataset to edit and choose *View Permissions*.



Select [here](#) and you will be directed to the **Edit ACL** page.



NFSv4 ACL Editor

Field	Description
Path	Shows the full pathway to the file.
Owner	User who controls the dataset. This user always has permissions to read or write the ACL and read or write attributes. Users created manually or imported from a directory service appear in the drop-down menu.
Owner Group	The group which controls the dataset. This group has the same permissions as granted to the group@ Who. Groups created manually or imported from a directory service appear in the drop-down menu.

You can select any user accounts or groups imported from a directory service as the primary *User* or *Group*.

Access Control List

To add a new item to the ACL, click *Add Item*, define *Who* the Access Control Entry (ACE) applies to, and configure permissions and inheritance flags for the ACE.

Field	Description
Add Item	Adds a new ACE to the Access Control List.
Who	Access Control Entry (ACE) user or group. Select a specific User or Group for this entry. Owner@ applies this entry to the user that owns the dataset. Group@ applies this entry to the group that owns the dataset. Everyone@ applies this entry to all users and groups. See nfs4_setfac(1) NFSv4 ACL ENTRIES .

ACL Type	How the Permissions apply to the chosen Who. Choose Allow to grant the specified permissions and Deny to restrict the specified permissions.
Permissions Type	Basic shows general permissions. Advanced shows each permission type for finer control.
Permissions	Select permissions to apply to the chosen Who. Choices change depending on the Permissions Type.
Flags Type	Select the set of ACE inheritance Flags to display. Basic shows nonspecific inheritance options. Advanced shows specific inheritance settings for finer control.
Flags	How this ACE applies to newly created directories and files within the dataset. Basic flags enable or disable ACE inheritance. Advanced flags allow further control of how the ACE applies to files and directories in the dataset.
Strip ACL	This action removes all ACLs from the current dataset and any directories or files contained within this dataset. Stripping the ACL resets dataset permissions and can make data inaccessible until you create new permissions.
Use ACL Preset	Choosing an entry loads a preset ACL configured to match general permissions situations. The chosen preset will REPLACE the ACL currently displayed in the form and delete any unsaved changes. The preset options are <i>NFS4_OPEN</i> , <i>NFS4_RESTRICTED</i> , or <i>NFS4_HOME</i> .

Permissions and Flags

TrueNAS divides permissions into Basic and Advanced options. The basic options are commonly-used groups of advanced options.

Basic inheritance flags only enable or disable ACE inheritance. Advanced flags offer finer control for applying an ACE to new files or directories.

Basic Permissions

Permission	Description
Read (r-x---a-R-c--)	View file or directory contents, attributes, named attributes, and ACL.
Modify (rwxpDdaARwC--s)	Adjust file or directory contents, attributes, and named attributes. Create new files or subdirectories. Includes the <i>Traverse</i> permission.
Traverse (--x---a-R-c--)	Execute a file or move through a directory.
Full Control (rwxpDdaARwCcos)	Apply all permissions.

Advanced Permissions

Permission	Description
Read Data (r)	View file contents or list directory contents.
Write Data (w)	Create new files or modify any part of a file.
Append Data (p)	Add new data to the end of a file.
Read Named Attributes (R)	View the named attributes directory.
Write Named Attributes (W)	Create a named attribute directory. Must be paired with the Read Named Attributes permission.
Execute (x)	Execute a file, move through, or search a directory.
Delete Children (D)	Delete files or subdirectories from inside a directory.
Read Attributes (a)	View file or directory non-ACL attributes.
Write Attributes (A)	Change file or directory non-ACL attributes.
Delete (d)	Remove the file or directory.
Read ACL (c)	View the ACL.
Write ACL (C)	Change the ACL and the ACL mode.
Write Owner (o)	Change the user and group owners of the file or directory.
Synchronize (s)	Synchronous file read/write with the server. This permission does not apply to FreeBSD clients.

Basic Flags

Flag	Description
Inherit (fd-----)	Enable ACE inheritance.
No Inherit (------)	Disable ACE inheritance.

Advanced Flags

Flag	Description
File Inherit (f)	The ACE is inherited with subdirectories and files. It applies to new files.
Directory Inherit (d)	New subdirectories inherit the full ACE.
No Propagate Inherit (n)	The ACE can only be inherited once.
Inherit Only (i)	Remove the ACE from permission checks but allow new files or subdirectories to inherit it. Inherit Only is removed from these new objects.
Inherited (I)	Set when this dataset inherits the ACE from another dataset.

4.1.4 - Encryption

- - [Pool Manager Encryption](#)
 - [Encrypting a New Dataset](#)
 - [Changing Dataset Encryption](#)
 - [Locking and Unlocking Datasets](#)
 - [Locking a Dataset](#)
 - [Unlocking a Dataset](#)
 - [Encrypting a Zvol](#)
 - [Managing Encryption Credentials](#)
 - [Unlocking a Replicated Encrypted Dataset or Zvol Without a Passphrase](#)

TrueNAS SCALE offers ZFS encryption for your sensitive data in pools and datasets or zvols.

Users are responsible for backing up and securing encryption keys and passphrases! Losing the ability to decrypt data is similar to a catastrophic data loss.

Data-at-rest encryption is available with:

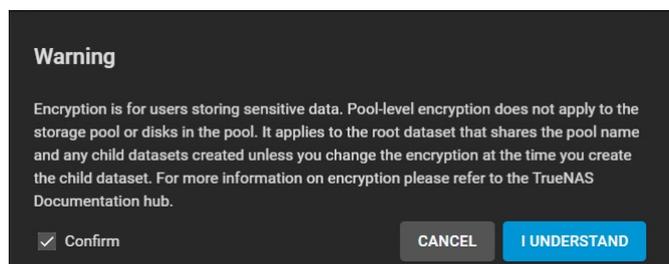
- [Self Encrypting Drives \(SEDs\)](#) using OPAL or FIPS 140.2 (Both [AES 256](#))
- Encryption of specific datasets (AES-256-GCM)

The local TrueNAS system manages keys for data-at-rest. Users are responsible for storing and securing their keys. TrueNAS SCALE includes the [Key Management Interface Protocol \(KMIP\)](#).

Pool Manager Encryption

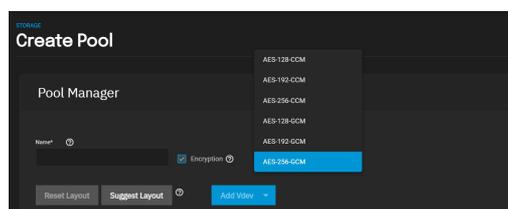
Encryption is for users storing sensitive data. Pool-level encryption does *NOT* apply to the storage pool or the disks in the pool. It only applies to the root dataset that shares the same name as the pool. Child datasets, or zvols, inherit encryption from the parent dataset unless you overwrite encryption when creating the child datasets or zvols.

Encrypting the root dataset of a new storage pool further increases data security. [Create a new pool](#) and check the **Encryption** box on the **Pool Manager** screen. The SCALE encryption warning dialog box displays.



Read the warning, select the **Confirm** checkbox, and click **I UNDERSTAND**.

You can select any of the encryption ciphers listed but we recommend using the default encryption cipher.



What are these options?

TrueNAS supports AES [Galois Counter Mode \(GCM\)](#) and [Counter with CBC-MAC \(CCM\)](#) algorithms for encryption. These algorithms provide authenticated encryption with block ciphers.

Encrypting a New Dataset

You can create new datasets within an existing storage pool as either encrypted or non-encrypted. A mix of encrypted and non-encrypted datasets can exist in a single storage pool.

To encrypt a dataset, [create a new dataset](#) and after typing a name scroll down to **Encryption Options**. The **Add Dataset** configuration screen encryption fields change based on the **Encryption Type** selected.

Inherit Checkbox

Because child datasets inherit settings from the parent dataset, the **Add Dataset** configuration screen displays with the inherit checkbox already checked. This means the inherit checkbox text for the child configuration screen changes based on the parent encryption setting.

Inherit (encrypted) displays for an encrypted parent dataset.

Add Dataset [Close]

Name and Options

Name * ⓘ

Comments ⓘ

Sync ⓘ

Compression level ⓘ

Enable Atime ⓘ

Encryption Options

Inherit (non-encrypted) ⓘ

Inherit (non-encrypted) displays for a parent dataset not encrypted.

Add Dataset [Close]

Name and Options

Name * ⓘ

Comments ⓘ

Sync ⓘ

Compression level ⓘ

Enable Atime ⓘ

Encryption Options

Inherit (encrypted) ⓘ

You can change the inherited encrypted/non-encrypted state by unchecking the inherit box. This displays the **Encryption** checkbox already check-marked.

Encryption Checkbox

Click the **Inherit (encrypted)** or **Inherited (non-encrypted)** checkbox with the checkmark to turn off inherited encryption settings. The **Encryption** checkbox displays already check-marked. You can now change this dataset's encryption settings.

If you uncheck the **Encryption** checkbox on the **Add Dataset** configuration screen, the encryption fields no longer display and the new child dataset is not encrypted.

Encryption Options fields change based on the **Encryption Type** selected. There are two options, **Key** or **Passphrase**. The default setting is **Key**.

Add Dataset [X]

Encryption Options

Inherit (non-encrypted) ⓘ

Encryption ⓘ

Encryption Type ⓘ

Key ▾

Generate Key ⓘ

Algorithm * ⓘ

AES-256-GCM ▾

The **Generate Key** checkbox defaults to check-marked. If you uncheck it, the **Key*** text field displays below it. Type the encryption key you want to use into this field.

Add Dataset [X]

Encryption Options

Inherit (non-encrypted) ⓘ

Encryption ⓘ

Encryption Type ⓘ

Key ▾

Generate Key ⓘ

Key * ⓘ

Algorithm * ⓘ

AES-256-GCM ▾

If you change the **Encryption Type** to **Passphrase** new passphrase fields display.

Add Dataset [X]

Encryption Options

Inherit (non-encrypted) ⓘ

Encryption ⓘ

Encryption Type ⓘ

Passphrase ▾

Passphrase * ⓘ

Confirm Passphrase * ⓘ

pbkdf2iters * ⓘ

350000 ⓘ

Algorithm * ⓘ

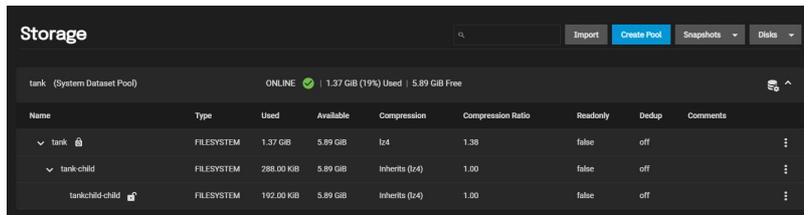
AES-256-GCM ▾

If using the passphrase option choose a complex phrase not easy to guess.

Keep both encryption keys and/or passphrases safeguarded in a secure and protected place. Losing encryption keys or passphrases can result in

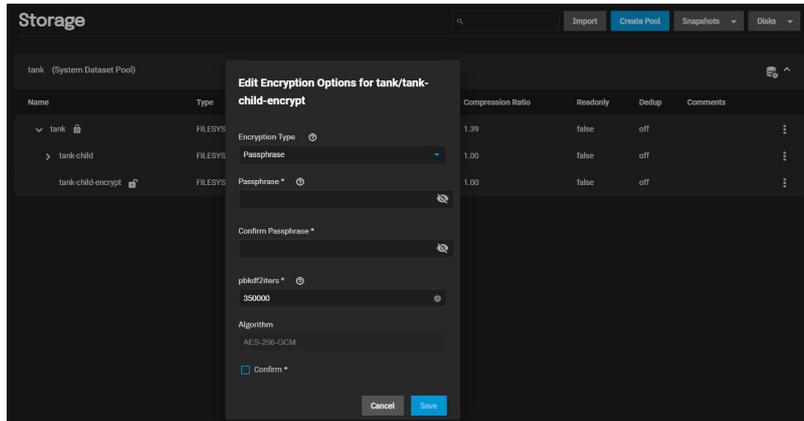
permanent data loss!

After configuring the new dataset encryption settings and any other settings, click **Save**. The new dataset displays on the **Storage** screen below its parent dataset. If you encrypt a dataset, an unlocked icon displays to the right of its name and a locked icon displays to the right of the root dataset name. A child dataset remain unlocked until you lock it.



Changing Dataset Encryption

Click on **Encryption Options** on the **Dataset Action** menu to change dataset encryption settings. This option only displays on the menu for datasets with encryption configured. The **Edit Encryption Options** configuration window displays and window name includes the dataset full path name. In the example used it includes the root dataset *tank*, the child dataset without encryption *tank-child*, and finally the selected child-of-the-child dataset with encryption *tank-child-encrypt* (i.e., *tank/tank-child/tank-child-encrypt*).



Click the **Confirm** checkbox to check-mark it and then click **Save** after making any changes.

Save any change to the encryption key or passphrase, and update your saved passcodes and keys file, and then back up that file.

Locking and Unlocking Datasets

TrueNAS displays a dataset status with icons:

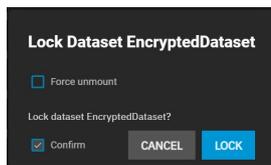
- Dataset unlocked icon:
- Dataset locked icon:

The locked icon displayed beside the root dataset after adding a dataset with encryption and also beside a dataset where the pool encryption properties don't match the root dataset is:

You can only lock and unlock an encrypted dataset when it is secured with a passphrase instead of a key file. Before locking a dataset, verify that it is not currently in use.

Locking a Dataset

Click the dataset's icon to display the **Dataset Actions** menu and then click on **Lock**. The **Lock Dataset** dialog box displays and includes the dataset full path name.

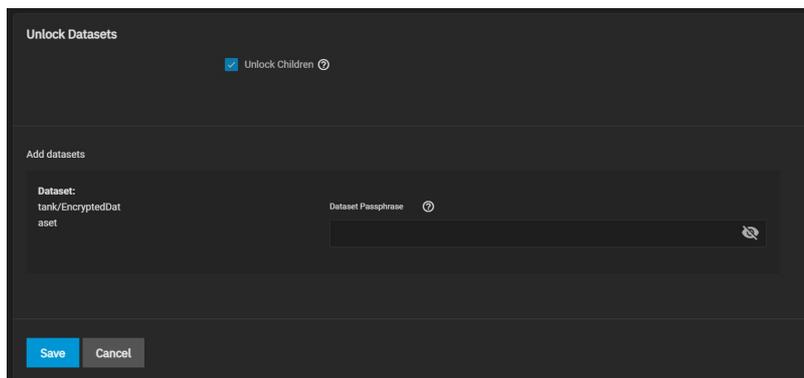


Use the **Force unmount** option only if you are certain no one is currently accessing the dataset. Click the **Confirm** checkbox to check-mark it and activate the **LOCK** button, and then click **LOCK**. A confirmation window displays indicating the dataset is locked and the unlock icon changes to a locked icon.

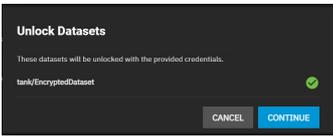
You *cannot* use locked datasets.

Unlocking a Dataset

To unlock a dataset, click on the icon to display the **Dataset Actions** menu and then click on **Unlock**.



Type the passphrase into the **Dataset Passphrase** field and click **Save**. You can unlock all locked child datasets using the same passphrase at the same time by check-marking the **Unlock Children** checkbox. A confirmation window displays.



Click **CONTINUE** to confirm you want to unlock the datasets or **CANCEL** to exit and keep the datasets locked. A second confirmation window displays confirming the datasets are unlocked. Click **CLOSE**. TrueNAS displays the dataset with the unlocked icon.

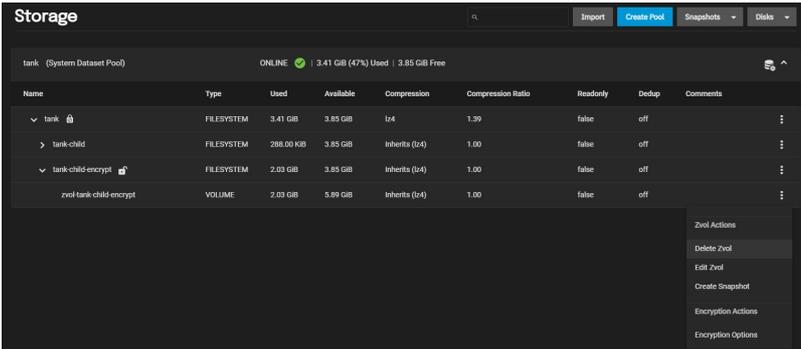
Encrypting a Zvol

Encryption is for securing sensitive data.

You can only encrypting a zvol if you create the zvol from a dataset with encryption.

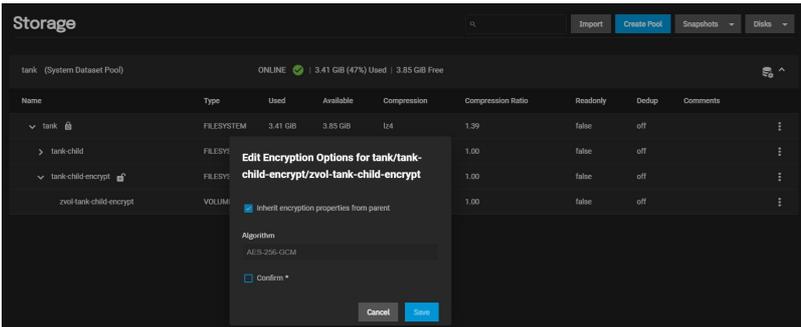
Users are responsible for backing up and securing encryption keys and passphrases! Losing the ability to decrypt data is similar to a catastrophic data loss.

Zvols, like datasets, inherit encryption settings from the parent dataset. To encrypt a zvol, select a dataset configured with encryption and then [create a new zvol](#). Next, click the  icon to display the **Zvol Actions** menu.



If you do not see encryption options on the menu you created the zvol from a dataset not configured with encryption. You can deleted the zvol and start over.

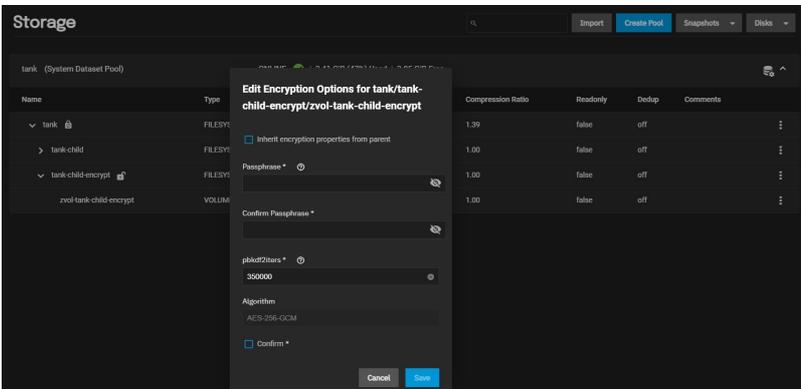
Click **Encryption Options**. The **Edit Encryption Options** configuration window displays with the **Inherit encryption properties from parent** checkbox already check-marked.



Like datasets, the window name includes the full path for the zvol. In this example, the root dataset *tank*, the encrypted child dataset *tank-child-encrypt*, and finally the zvol name *zvol-tank-child-encrypt* (i.e., *tank/tank-child-encrypt/zvol-tank-child-encrypt*).

If not making changes, click the **Confirm** checkbox to activate the **Save** button, and then click **Save**. The zvol is encrypted with settings inherited from its parent.

To change inherited encryption properties, click on the inherit checkbox to uncheck it. Additional configuration option fields display.



If **Encryption Type** is set to **Key**, type an encryption key into the **Key** field or check-mark the **Generate Key** checkbox. If set to **Passphrase**, type a passphrase at least eight characters long into both the **Passphrase** and **Confirm Passphrase** fields. After making any changes, click the **Confirm** checkbox to check-mark it and activate the **Save** button, and then click **Save**. The zvol is now encrypted with settings not inherited from its parent.

Save any change to the encryption key or passphrase, update your saved passcodes and keys file, and back up the file.

Managing Encryption Credentials

There are two ways to manage the encryption credentials, with a key file or passphrase.

Key Files

Creating a new encrypted pool automatically generates a new key file and prompts users to download it.

Always back up the key file to a safe and secure location.

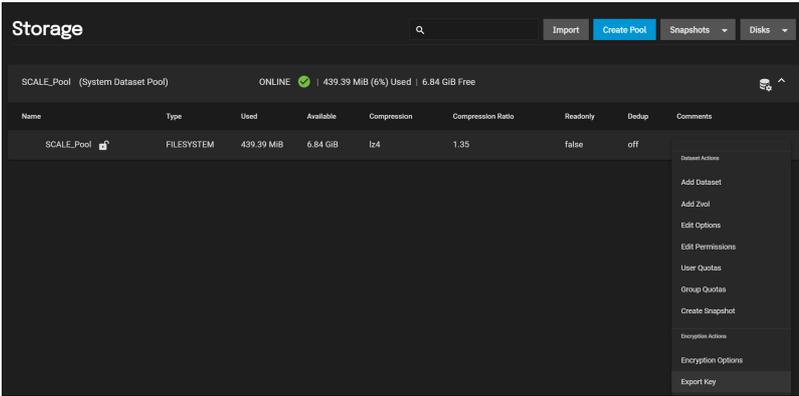
WARNING!

Losing the ability to unlock the pool can result in losing all data on the disks with no chance of recovery. Always back up the encryption key file or passphrase for an encrypted pool! The key file for an encrypted pool is secured in the system database and can be exported at any time from the pool options.

[Download Encryption Key](#) [Done](#)

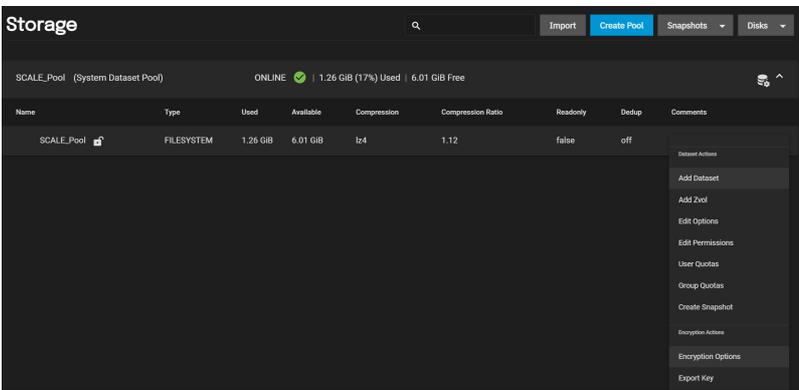
To manually back up a root dataset key file, open the pool menu and select **Export Key**. Next type the root user password to authorize the export.

menu and select **Export Key**. Next type the root user password to

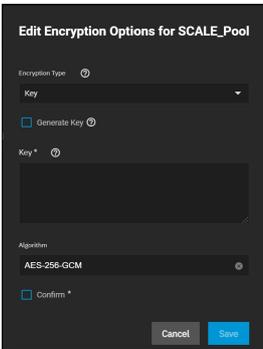


To change the key, click the dataset's

icon and then click on **Encryption Options**.



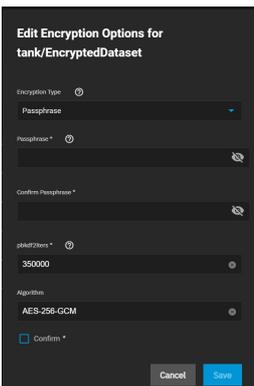
To enter your custom key click the **Generate Key** checkbox to uncheck it and display the **Key** text entry field. Leave the **Generate Key** check-marked to generate a random encryption key that displays in the **Key** field. Click **Save** to complete the process and close the window.



Passphrases

To use a passphrase instead of a key file, click the dataset's **Encryption Options**. Change the **Encryption Type** from **Key** to **Passphrase**.

icon to display the **Dataset Actions** menu and then click on



Set the rest of the options:

- **Passphrase:** A user-defined string at least eight characters long that is required to decrypt the dataset. Type it into the **Passphrase** and **Confirm Passphrase** fields.

The passphrase is the only means to decrypt the information stored in this dataset.

Be sure to create a memorable passphrase or physically secure the passphrase.

- **pbkdf2iters**: The number of password-based key derivation function 2 ([PBKDF2](#)) iterations to use for reducing vulnerability to brute-force attacks. Users must enter a number greater than 100000.

Unlocking a Replicated Encrypted Dataset or Zvol Without a Passphrase

TrueNAS SCALE users should either replicate the dataset/Zvol without properties to disable encryption at the remote end or construct a special json manifest to unlock each child dataset/zvol with a unique key.

Method 1: Construct JSON Manifest

1. Replicate every encrypted dataset you want to replicate with properties.
2. Export key for every child dataset that has a unique key.
3. For each child dataset construct a proper json with poolname/datasetname of the destination system and key from the source system like this:

```
{"tank/share01": "57112db4be777d93fa7b76138a68b790d46d6858569bf9d13e32eb9fda72146b"}
```
4. Save this file with the extension .json.
5. On the remote system, unlock the dataset(s) using properly constructed json files.

Method 2: Replicate Encrypted Dataset/zvol Without Properties

Uncheck properties when replicating so that the destination dataset will not be encrypted on the remote side and will not require a key to unlock.

1. Go to **Data Protection** and click **ADD** in the *Replication Tasks* window.
2. Click *Advanced Replication Creation*.
3. Fill out the form as needed and make sure *Include Dataset Properties* is **NOT** checked.
4. Click **Save**.

Method 3: Replicate Key Encrypted Dataset/zvol

Check **Full Filesystem Replication** so that the destination dataset will use the exported Encryption key from the source pool/dataset to unlock.

1. Go to **Storage -> pool/root dataset**. Click **!** and select **Export Key**.
2. Download the key, open the text file, and copy the Key code.
3. Go to **Data Protection** and click **ADD** in the *Replication Tasks* window.
4. Click *Advanced Replication Creation*.
5. Fill out the form as needed and make sure to enable *Full Filesystem Replication*.
6. Click **Save**.
7. On the receiving pool/dataset:
 - Click **!** next to pool/dataset and select *Unlock*.
 - Unset *Unlock with Key file*.
 - Paste the Key Code into Dataset Key. (if there is a space character at the end of the key - delete the space.)
 - Click **Save**.
 - Click *Continue*.

4.1.5 - Fusion Pools

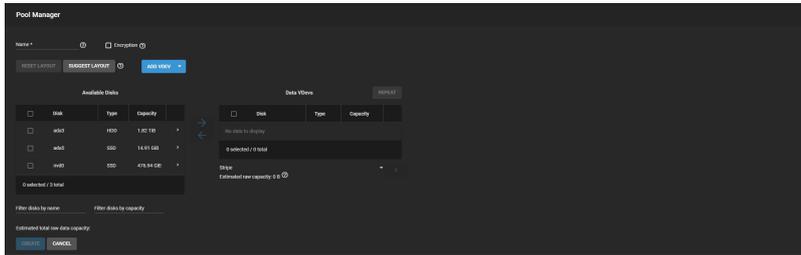
Fusion Pools are also known as **ZFS Allocation Classes**, **ZFS Special vdevs**, and **Metadata vdevs**.

What's a special vdev?

A special vdev can store meta data such as file locations and allocation tables. The allocations in the special class are dedicated to specific block types. By default, this includes all metadata, the indirect blocks of user data, and any deduplication tables. The class can also be provisioned to accept small file blocks. This is a great use case for high performance but smaller sized solid-state storage. Using a special vdev drastically speeds up random I/O and cuts the average spinning-disk I/Os needed to find and access a file by up to half.

Creating a Fusion Pool

Go to **Storage > Pools**, click **ADD**, and select *Create new pool*.



A pool must always have one normal (non-dedup/special) vdev before other devices can be assigned to the special class. Configure the **Data VDevs**, then click **ADD VDEV** and select *Metadata*.

Add SSDs to the new **Metadata VDev** and select the same layout as the **Data VDevs**.

The metadata special vdev is critical for pool operation and data integrity, so you must protect it with hot spare(s).

UPS Recommendation

When using SSDs with an internal cache, add Uninterruptible Power Supply (UPS) to the system to help minimize the risk from power loss.

Using special vdevs identical to the data vdevs (so they can use the same hot spares) is recommended, but for performance reasons you can make a different type of vdev (like a mirror of SSDs). In that case you must provide hot spare(s) for that drive type as well. Otherwise, if the special vdev fails and there is no redundancy, the pool becomes corrupted and prevents access to stored data.

Drives added to a metadata vdev cannot be removed from the pool.

When more than one metadata vdev is created, then allocations are load-balanced between all these devices. If the special class becomes full, then allocations spill back into the normal class.

After the fusion pool is created, the **Status** shows a **Special** section with the metadata SSDs.

See [Pool Operations](#) for more information on managing pools.

4.2 - Snapshots

- [Snapshot Creation Options](#)

Snapshots are one of the most powerful features of ZFS. A snapshot provides a read only point-in-time copy of a file system or volume. This copy does not consume extra space in the ZFS pool. The snapshot only records the differences between storage block references whenever the data is modified.

Why do I want to keep snapshots?

Snapshots keep a history of files and provide a way to recover an older or even deleted files. For this reason, many administrators take regular snapshots, store them for some time, and copy them to a different system. This strategy allows an administrator to roll the system data back to a specific point in time. In the event of catastrophic system or disk failure, off-site snapshots can restore data up to the most recent snapshot.

Taking snapshots requires the system have all [pools](#), [datasets](#), and [zvols](#) already configured.

Snapshot Creation Options

Single Snapshots

Creating a Single Snapshot

Consider making a [Periodic Snapshot Task](#) to save time and create regular, fresh snapshots.

Video Tutorial

This short video demonstrates manually adding a snapshot



Video Player is loading.
Video URL: <https://www.ixenas.com/docs/files/scaleangelfishmanualsnapshots.mp4>

Play Video

Go to **Storage** and click **Snapshots**, then click **ADD**.

Snapshot

Dataset *

Name

manual-2021-08-31_07-22

Naming Schema

Recursive

Save Cancel

Use the **Dataset** dropdown to select an existing ZFS pool, dataset, or zvol to snapshot.

TrueNAS software generates a suggested name that you can override with any custom string.

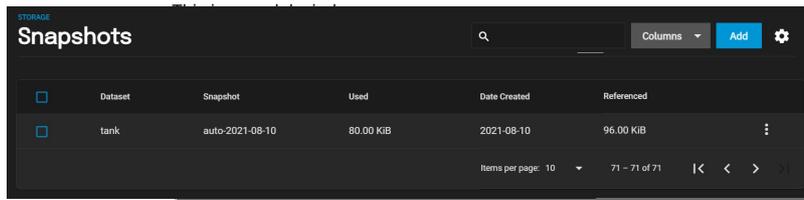
TrueNAS software populates the **Naming Schema** drop-down with periodic snapshot task schemas already created. Choosing one generates a name for the snapshot using the naming schema. If you select Periodic Snapshot and replicates that snapshot. You cannot enter a value in **Naming Schema** and **Name** as selecting or entering a value in **Naming Schema** populates the other.

To include child datasets with the snapshot, set **Recursive**.

- Unknown, selected

Managing Snapshots

Go to **Storage** and click **Snapshots** to manage created snapshots.



Each entry in the list includes the dataset and snapshot names. Entries also display the snapshot numbers, the space they use, the date the system created them, and the amount of data the dataset can access.

Click to view snapshot options.

Delete

The **Delete** option destroys the snapshot. You must delete child clones before you can delete their parent snapshot. While creating a snapshot is instantaneous, deleting one is I/O intensive and can take a long time, especially when deduplication is enabled.

Why? ZFS has to review all allocated blocks before deletion to see if another process is using that block. If not used, the ZFS can free that block.

Clone to New Dataset

The **Clone to New Dataset** option creates a new snapshot *clone* (dataset) from the snapshot contents.

What is a clone?
A **clone** is a writable copy of the snapshot. Because a clone is a mountable dataset, it appears in the **Storage** screen rather than the **Snapshots** screen. By default, TrueNAS adds **-clone** to the new snapshot name when creating the clone.

A dialog prompts for the new dataset name. The suggested name derives from the snapshot name.

Rollback

The **Rollback** option reverts the dataset back to the point in time saved by the snapshot.

Rollback is a dangerous operation that causes any configured replication tasks to fail. Replications use the existing snapshot when doing an incremental backup, and rolling back can put the snapshots 'out of order'. To restore the data within a snapshot, the recommended steps are:

1. Clone the desired snapshot.
2. Share the clone with the share type or service running on the TrueNAS system.
3. Allow users to recover their needed data.
4. Delete the clone from **Storage**.

This approach does not destroy any on-disk data or impact replication.

TrueNAS asks for confirmation before rolling back to the chosen snapshot state. Clicking **Yes** reverts all dataset files to the state they were in when TrueNAS created the snapshot.

Batch Operations

To delete multiple snapshots, select the left column box for each snapshot to include. Click the **Delete** button that displays.

To search through the snapshots list by name, type a matching criteria into the **Filter Snapshots** text field. The list now displays only the snapshot names that match the filter text.

Browsing a Snapshot Collection

Browsing a snapshot collection is an advanced capability that requires ZFS and command-line experience.

All dataset snapshots are accessible as an ordinary hierarchical file system, accessed from a hidden `.zfs` located at the root of every dataset.

A snapshot and any files it contains are not accessible or searchable if the snapshot mount path is longer than 88 characters. The data within the snapshot is safe but to make the snapshot accessible again shorten the mount path.

A user with permission to access the hidden file can view and explore all snapshots for a dataset from the **Shell** or the **Shares** screen using services like **SMB**, **NFS**, and **SFTP**.

In summary, the main required changes to settings are:

- In dataset properties, change the ZFS properties to enable snapshot visibility.
- In the Samba auxiliary settings, change the `veto files` command to not hide the `.zfs`, and add the setting `zfsacl:expose_snapdir=true`.

The effect is that any user who can access the dataset contents can view the list of snapshots by going to the dataset `.zfs` directory. Users can browse and search any files they have permission to access throughout the entire dataset snapshot collection.

When creating a snapshot, permissions or ACLs set on files within that snapshot might limit access to the files.

Snapshots are read-only, so users do not have permission to modify a snapshot or its files, even if they had write permissions when creating the snapshot.

The ZFS `zfs diff` command, which can run in the **Shell**, lists all changed files between any two snapshot versions within a dataset, or between any snapshot and the current data.

VMware-Snapshots

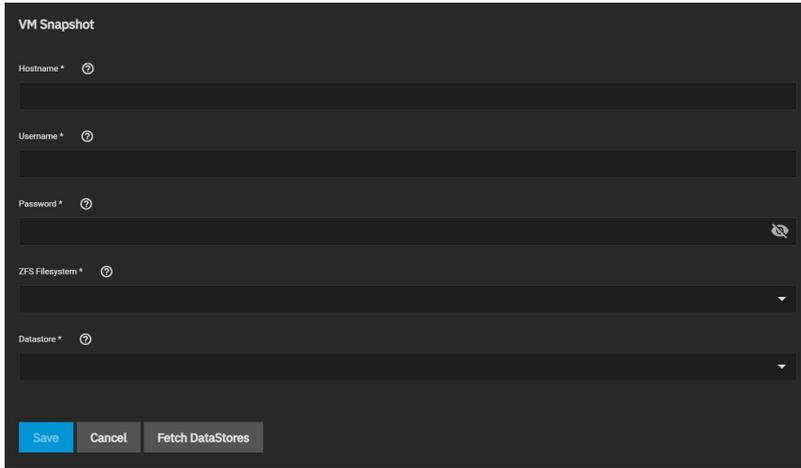
VMware-Snapshots coordinate ZFS snapshots when using TrueNAS as a VMware datastore. When TrueNAS creates a ZFS snapshot, it snapshots any running VMware virtual machines before it takes a scheduled or manual ZFS snapshot of the dataset or zvol backing that VMware datastore.

You must power on virtual machines for TrueNAS to copy snapshots to VMware. The temporary VMware snapshots deleted on the VMware side, still exist in the ZFS snapshot and are available as stable restore points. These coordinated snapshots go in the **Snapshots** list.

Only paid versions of VMware ESXi support VMware-Snapshots. Attempting to create VMware-Snapshots with ESXi free results in the following error message: **"Error: Can't create snapshot, current license or ESXi version prohibits execution of the requested operation."** ESXi free has a locked (read only) API that prohibits using TrueNAS VMware-Snapshots. *VMware vSphere Essentials Kit* is the cheapest ESXi edition that is compatible with TrueNAS VMware-Snapshots.

Create a VMware Snapshot

Go to **Storage** and click **VMware Snapshots**, then click **ADD**.



VM Snapshot

Hostname *

Username *

Password *

ZFS Filesystem *

Datastore *

Save Cancel Fetch DataStores

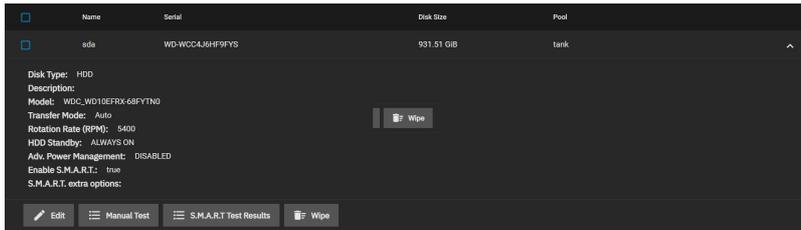
Setting	Description
Hostname	Enter the IP address or hostname of the VMware host. When clustering, this is the vCenter server for the cluster.
Username	Enter the user on the VMware host with permission to snapshot virtual machines.
Password	Enter the password associated with the user entered in Username .
ZFS Filesystem	Select a file system to snapshot.
Datastore	After entering the values in Hostname , *Username , and Password , click Fetch DataStores and select the datastore to synchronize.

TrueNAS connects to the VMware host after clicking **Fetch DataStores**. The **ZFS Filesystem** and **Datastore** drop-down menus populate from the VMware host response. Choosing a datastore also selects any previously mapped dataset.

4.3 - Disks

The *Disks* page displays the names, serial numbers, sizes, and pools of all the system's physical drives. Users can customize disk columns using the *Columns* drop-down*.

Clicking the  in a disk's row will expand it to show the traits specific to that disk.



Managing Disks

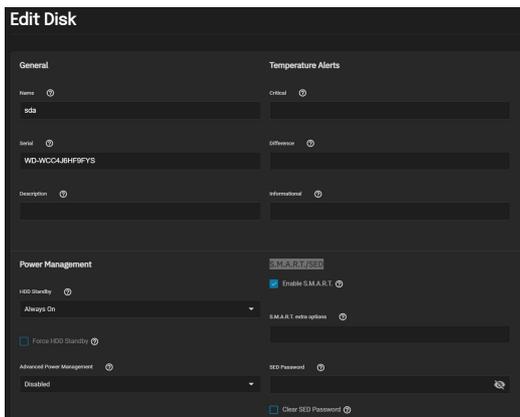
Managing Disks

To manage disks, go to **Storage** and click *Disks*, then select *Disks*.

The *Disks* page lets users edit disks, perform manual tests, and view S.M.A.R.T. test results. Users may also delete obsolete data off an unused disk.

Editing Disks

Clicking *Edit* allows users to configure general disk settings, as well as power management, temperature alerts, and S.M.A.R.T./SED settings.



General

Setting	Description
Name	Linux disk device name.
Serial	Serial number for this disk.
Description	Notes about this disk.

Power Management

Setting	Description
HDD Standby	Minutes of inactivity before the drive enters standby mode. This forum post describes identifying spun down drives. Temperature monitoring is disabled for standby disks.
Force HDD Standby	Allows the drive to enter standby, even when non-physical S.M.A.R.T. operations could prevent the drive from sleeping.
Advanced Power Management	Select a power management profile from the menu.

Temperature Alerts

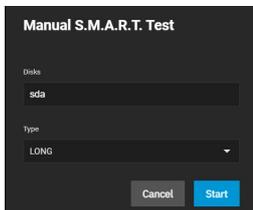
Setting	Description
Critical	Threshold temperature in Celsius. If the drive temperature is higher than this value, a LOG_CRIT level log entry is created and an email is sent. 0 disables this check.
Difference	Report if the temperature of a drive has changed by this many degrees Celsius since the last report. 0 disables the report.
Informational	Report if drive temperature is at or above this temperature in Celsius. 0 disables the report.

S.M.A.R.T./SED

Setting	Description
Enable S.M.A.R.T.	Enabling allows the system to conduct periodic S.M.A.R.T. tests .
S.M.A.R.T. extra options	Additional smartctl(8) options.
SED Password	Set or change the password of this SED. This password is used instead of the global SED password.
Clear SED Password	Clear the SED password for this disk.

Manual Testing

Select the disk(s) you want to perform a S.M.A.R.T. test on and click *Manual Test*.



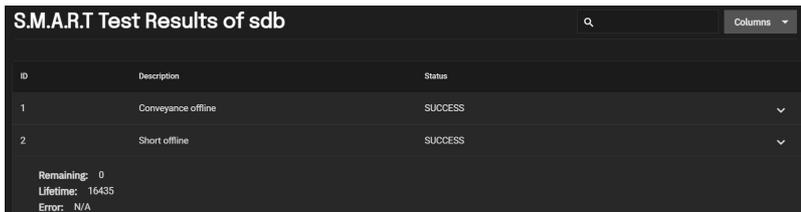
- *Long* - runs SMART Extended Self Test. This will scan the entire disk surface and can take many hours on large-volume disks.
- *Short* - runs SMART Short Self Test (usually under ten minutes). These are basic disk tests that vary by manufacturer.
- *Conveyance* - runs a SMART Conveyance Self Test. This self-test routine is intended to identify damage incurred during transporting of the device. This self-test routine requires only minutes to complete.
- *Offline* - runs SMART Immediate Offline Test. The effects of this test are visible only in that it updates the SMART Attribute values, and if the test finds errors, they appear in the SMART error log.

Click *Start* to begin the test. Depending on the test type you choose, the test can take some time to complete. TrueNAS generates alerts when tests discover issues.

For information on automated S.M.A.R.T. testing, see the [S.M.A.R.T. tests](#) SCALE article.

S.M.A.R.T. Test Results

To review test results, expand the disk and click *S.M.A.R.T. Test Results*.



Users can also view S.M.A.R.T. Test Results in **Shell** using `smartctl` and the name of the drive: `smartctl -l selftest /dev/sdb`.

Wipe

The *Wipe* option deletes obsolete data off an unused disk.

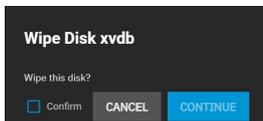
Wipe is a destructive action and results in permanent data loss! Back up any critical data before wiping a disk.

TrueNAS only shows the *Wipe* option for unused disks. Click *Wipe* to open a dialog with additional options:

- *Quick* - Erases only the partitioning information on a disk without clearing other old data, making it easy to reuse. Quick wipes take only a few seconds.
- *Full with zeros* - Overwrites the entire disk with zeros and can take several hours to complete.
- *Full with random* - Overwrites the entire disk with random binary code and takes even longer than *Full with zeros* to complete.

Ensure you have backed-up all data and are no longer using the disk. Triple check that you have selected the correct disk for the wipe. Recovering data from a wiped disk is usually impossible.

After choosing the appropriate method, click *Wipe* and confirm the action.



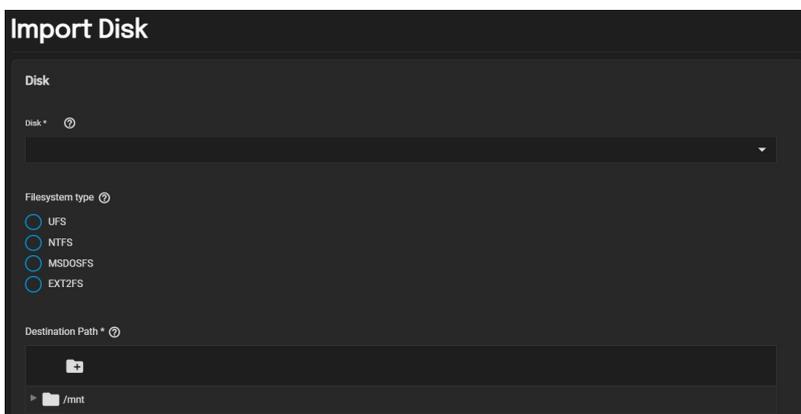
Verify the name to ensure you have chosen the correct disk. When satisfied the disk can be wiped, set *Confirm* and click *Continue*.

Importing Disks

Go to **Storage** and click *Disks*, then select *Import Disk* to integrate UFS (BSD Unix), NTFS (Windows), MSDOS (FAT), or EXT2 (Linux) formatted disks into TrueNAS. Importing is a one-time procedure that copies the data from that disk into a TrueNAS dataset. TrueNAS can only import one disk at a time, and it must be installed or physically connected to the TrueNAS system.

What about EXT3 or EXT4 filesystems?

Importing an EXT3 or EXT4 filesystem is possible in some cases, although neither is fully supported. EXT3 journaling is not supported, so those filesystems must have an external `fsck` utility, like the one provided by [E2fsprogs utilities](#), run on them before import. EXT4 filesystems with extended attributes or inodes greater than 128 bytes are not supported. EXT4 filesystems with EXT3 journaling must have an `fsck` run on them before import, as described above.



Use the drop-down menu to select the *Disk* to import.

TrueNAS attempts to detect and select the *Filesystem type*. Selecting the *MSDOSFS* filesystem shows an additional *MSDOSFS locale* drop-down menu. Use this option to select the locale when non-ASCII characters are present on the disk.

Finally, select the ZFS dataset you want to hold the copied data in *Destination Path*.

After clicking *Save*, the chosen *Disk* mounts and copies its contents to the specified dataset at the end of the *Destination Path*. To monitor an in-progress import, open the Task Manager by clicking the  in the interface top bar. The disk unmounts after the copy operation completes. A dialog allows viewing or downloading the disk import log.

The import was interrupted!

Use the same import procedure to restart the task. Choose the same *Destination Path* as the interrupted import for TrueNAS to scan the destination for previously imported files and resume importing any remaining files.

Replacing Disks

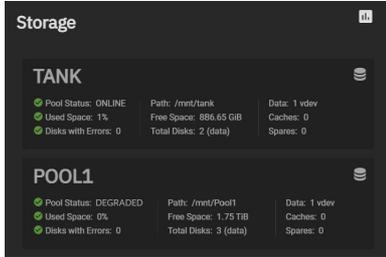
Hard drives and solid-state drives (SSDs) have a finite lifetime and can fail unexpectedly. When a disk fails in a Stripe (RAID0) pool, must to recreate the entire pool and restore all data backups. We always recommend creating non-stripe storage pools that have disk redundancy.

To prevent further redundancy loss or eventual data loss, always replace a failed disk as soon as possible! TrueNAS integrates new disks into a pool to restore it to full functionality.

Replacing a Disk

TrueNAS requires another disk of the same or greater capacity to replace a failed disk. The disk must be installed in the TrueNAS system and not part of an existing storage pool. TrueNAS wipes any data on the replacement disk as part of the process.

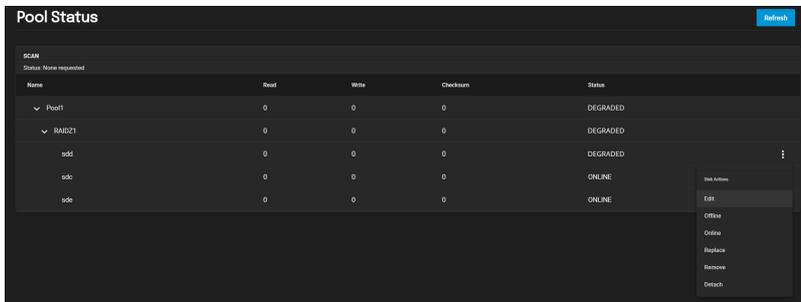
The TrueNAS **Dashboard** shows when a disk failure degrades a pool.



Click the  on the pool card to go to the *Pool Status* screen and locate the failed disk.

Offline the Failed Disk

Clicking  next to the failed disk shows additional operations.



We recommend users *Offline* the disk before starting the replacement. Doing so removes the device from the pool and can prevent swap issues.

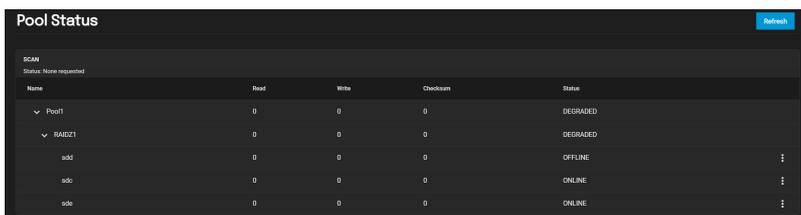
Can I use a disk that is failing but still active?

There are situations where a disk that has not completely failed can be left online to provide additional redundancy during the replacement procedure. **We do not recommend leaving failed disks online unless you know the exact condition of the failing disk.** Attempting to replace a heavily degraded disk without offlineing it will significantly slow down the replacement process.

The offline failed?

If the *Offline* operation fails with a "Disk offline failed - no valid replicas" message, go to **Storage**, click the  for the degraded pool, and select *Scrub Pool*. When the scrub operation finishes, reopen the pool *Status* and try to *Offline* the disk again.

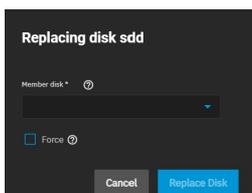
You can physically remove the disk from the system when the disk status is *Offline*.



If the replacement disk is not already physically added to the system, add it now.

Online the New Disk

In *Pool Status*, open the options for the *Offline* disk and click *Replace*.



Select a new member disk and click *Replace Disk*. The new disk must have the same or greater capacity as the disk you are replacing. The replacement fails when the chosen disk has partitions or data present. To **destroy** any data on the replacement disk and allow the replacement to continue, set the *Force* option.

When the disk wipe completes, TrueNAS will start replacing the failed disk. *Pool Status* updates to show the in-progress replacement.

Pool Status Refresh

RESILVER
 Status: SCANNING
 Completed: 47.965
 Time Remaining: 7 seconds
 Errors: 0
 Date: 2021-09-01 13:28:14

Name	Read	Write	Checksum	Status
Pool1	0	0	0	ONLINE
RAIDZ1	0	0	0	ONLINE
REPLACING	0	0	0	ONLINE

TrueNAS resilvers the pool during the replacement process. For pools with large amounts of data, this can take a long time. When the resilver is complete, the pool status returns to **Online** shows the new disk.

Pool Status Refresh

RESILVER
 Status: FINISHED
 Errors: 0
 Date: 2021-09-01 13:28:14

Name	Read	Write	Checksum	Status
Pool1	0	0	0	ONLINE
RAIDZ1	0	0	0	ONLINE
sdb	0	0	0	ONLINE
sdc	0	0	0	ONLINE

5 - Shares

File sharing is one of the primary benefits of a NAS. TrueNAS helps foster collaboration between users through network shares. TrueNAS SCALE allows users to create and configure block (iSCSI) shares targets, Windows SMB shares, Unix (NFS) shares, and WebDAV shares.

When creating zvols for shares, avoid giving them names with capital letters or spaces since they can cause problems and failures with iSCSI and NFS shares.



Block (iSCSI) Shares Targets

iSCSI (Internet Small Computer Systems Interface) represents standards for using Internet-based protocols for linking binary data storage device aggregations. IBM and Cisco submitted the draft standards in March 2000. Since then, iSCSI has seen widespread adoption into enterprise IT environments.

iSCSI functions through encapsulation. The OSI (Open Systems Interconnection Model) encapsulates SCSI commands and storage data within the session stack. The OSI further encapsulates the session stack within the transport stack, the transport stack within the network stack, and the network stack within the data stack. Transmitting data this way permits block-level access to storage devices over LANs, WANs, and even the Internet itself (although performance may suffer if your data traffic is traversing the Internet).

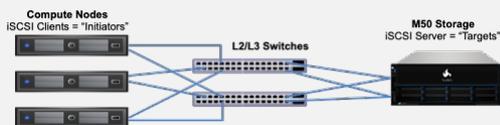
The table below shows where iSCSI sits in the OSI network stack:

OSI Layer Number	OSI Layer Name	Activity as it relates to iSCSI
7	Application	An application tells the CPU that it needs to write data to non-volatile storage.
6	Presentation	OSI creates a SCSI Command, SCSI Response, or SCSI data payload to hold the application data and communicate it to non-volatile storage.
5	Session	Communication between the source and the destination devices begins. This communication establishes when the conversation starts, what it will talk about, and when the conversation ends. This entire dialogue represents the session. OSI encapsulates the SCSI Command, SCSI Response, or SCSI data payload containing the application data within an iSCSI Protocol Data Unit (PDU).
4	Transport	OSI encapsulates the iSCSI PDU within a TCP segment.
3	Network	OSI encapsulates the TCP segment within an IP packet.
2	Data	OSI encapsulates the IP packet within the Ethernet frame.
1	Physical	The Ethernet frame transmits as bits (zeros and ones).

Unlike other sharing protocols on TrueNAS, an iSCSI share allows block sharing *and* file sharing. Block sharing provides the benefit of [block-level access](#) to data on the TrueNAS. iSCSI exports disk devices (zvols on TrueNAS) over a network that other iSCSI clients (initiators) can attach and mount.

iSCSI Terminology

- **CHAP (Challenge-Handshake Authentication Protocol):** an authentication method that uses a shared secret and three-way authentication to determine if a system is authorized to access the storage device. It also periodically confirms that the session has not been hijacked by another system. In iSCSI, the client (initiator) performs the CHAP authentication.
- **Mutual CHAP:** a CHAP type in which both ends of the communication authenticate to each other.
- **Internet Storage Name Service (iSNS):** protocol for the automated discovery of iSCSI devices on a TCP/IP network.
- **Extent:** the storage unit to be shared. It can either be a file or a device.
- **Portal:** indicates which IP addresses and ports to listen on for connection requests.
- **Initiators and Targets:** iSCSI introduces the concept of *initiators* and *targets* which act as sources and destinations respectively. iSCSI initiators and targets follow a client/server model. Below is a diagram of a typical iSCSI network. The TrueNAS storage array acts as the iSCSI target and can be accessed by many of the different iSCSI initiator types, including software and hardware-accelerated initiators.



The iSCSI protocol standards require that iSCSI initiators and targets be represented as iSCSI nodes. It also requires that each node be given a unique iSCSI name. To represent these unique nodes via their names, iSCSI requires the use of one of two naming conventions and formats, IQN or EUI. iSCSI also allows the use of iSCSI aliases which are not required to be unique and can be help manage nodes.

- **LUN: Logical Unit Number** representing a logical SCSI device. An initiator negotiates with a target to establish connectivity to a LUN. The result is an iSCSI connection that emulates a connection to a SCSI hard disk. Initiators treat iSCSI LUNs as if they were a raw SCSI or SATA hard drive. Rather than mounting remote directories, initiators format and directly manage filesystems on iSCSI LUNs. When configuring multiple iSCSI LUNs, create a new target for each LUN. Since iSCSI multiplexes a target with multiple LUNs over the same TCP connection, there can be TCP contention when more than one target accesses the same LUN. TrueNAS supports up to 1024 LUNs.
- **Jumbo Frames:** Jumbo frames are the name given to Ethernet frames that exceed the default 1500 byte size. This parameter is typically referenced by the nomenclature as maximum transmission unit (MTU). MTU that exceeds the default 1500 bytes necessitates that all devices transmitting Ethernet frames between the source and destination support the specific jumbo frame MTU setting, which means that NICs, dependent hardware iSCSI, independent hardware iSCSI cards, ingress and egress Ethernet switch ports, and the NICs of the storage array must all support the same jumbo frame MTU value. So, how does one decide if they should use jumbo frames?

Administrative time is consumed configuring Jumbo Frames and troubleshooting if/when things go sideways. Some network switches might also have ASICs optimized for processing MTU 1500 frames while others might be optimized for larger frames. Systems administrators should also account for the impact on host CPU utilization. Although Jumbo Frames are designed to increase data throughput, it may measurably increase latency (as is the case with some un-optimized switch ASICs); latency is typically more important than throughput in a VMware environment. Some iSCSI applications might see a net benefit running jumbo frames despite possible increased latency. Systems administrators should test jumbo frames on their workload with lab infrastructure as much as possible before updating the MTU on their production network.

TrueNAS Enterprise Feature:

- **ALUA: Asymmetric Logical Unit Access** allows a client computer to discover the best path to the storage on a TrueNAS system. HA storage clusters can provide multiple paths to the same storage. For example, the disks are directly connected to the primary computer and provide high speed and bandwidth when accessed through that primary computer. The same disks are also available through the secondary computer, but speed and bandwidth are restricted. With ALUA, clients automatically ask for and use the best path to the storage. If one of the TrueNAS HA computers becomes inaccessible, the clients automatically switch to the next best alternate path to the storage. When a better path becomes available, as when the primary host becomes available again, the clients automatically switch back to that better path to the storage.

Do not enable ALUA on TrueNAS unless it is also supported by and enabled on the client computers. ALUA only works when enabled on both the client and server.

iSCSI Configuration Methods

There are a few different approaches for configuring and managing iSCSI-shared data:

- TrueNAS CORE web interface: the TrueNAS web interface is fully capable of configuring iSCSI shares. This requires creating and populating [zvol block devices](#) with data, then setting up the [iSCSI Share](#). TrueNAS Enterprise licensed customers also have additional options to configure the share with [Fibre Channel](#).
- TrueNAS SCALE web interface: TrueNAS SCALE offers a similar experience to TrueNAS CORE for managing data with iSCSI; create and populate the block storage, then configure the iSCSI share.
- TrueCommand instances that have many TrueNAS systems connected can [manage iSCSI Volumes](#) from the TrueCommand web interface. TrueCommand allows creating block devices and configuring iSCSI Targets and Initiators from one central location.
- TrueNAS Enterprise customers that use vCenter to manage their systems can use the [TrueNAS vCenter Plugin](#) to connect their TrueNAS systems to vCenter and create and share iSCSI datastores. This is all managed through the vCenter web interface.

To get started with iSCSI shares, make sure you have already created a [zvol](#) or a [dataset](#) with at least one file to share.

Go to **Shares** and click **Configure** in the **Block (iSCSI) Shares Targets** window. You can either use the creation wizard or set one up manually.

Configuring an iSCSI Share Tutorial Video

This short tutorial video demonstrates basic steps to set up an iSCSI share configuration.



Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaleangelfishiscsi.mp4>

Play Video

Play

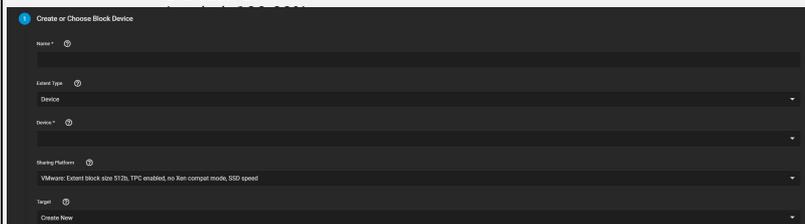
Mute

Current Time 0:00

Duration 1:44

Wizard Setup

Block Device



- Chapters

First, enter a name. It can only contain lowercase alphanumeric characters plus a dot (.), dash (-), or colon (:). We recommend keeping it short or at most 63 characters. Next, choose the **Extent Type**.

- If the **Extent Type** is **Device**, select the zvol to share from the **Device** menu.
- If the **Extent Type** is **File**, select the path to it and indicate the size.

Select the type of platform using the share. For example, if you use an updated Linux OS, choose **Modern OS**.

- captions off, selected

Portal

Audio Track

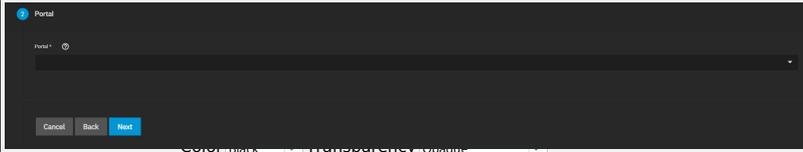
Now you either create a new portal or select an existing one from the dropdown list.

- Unknown, selected

If you create a new portal, you need to select a **Discovery Authentication Method**.

Fullscreen

If you set the **Discovery Authentication Method** to **CHAP** or **MUTUAL CHAP**, you also need to select a **Discovery Authentication Group**. If no group exists, click **Create New** from the dropdown list and enter a value in **Group ID**, **User**, and **Secret**.



When the **Discovery Authentication Method** is **NONE**, you can leave the **Discovery Authentication Group** empty.

Select **0.0.0.0** or **::** for the **IP Address** dropdown list, and click **NEXT**. **0.0.0.0** listens on all IPv4 addresses and **::** listens on all IPv6 addresses.

Font Size

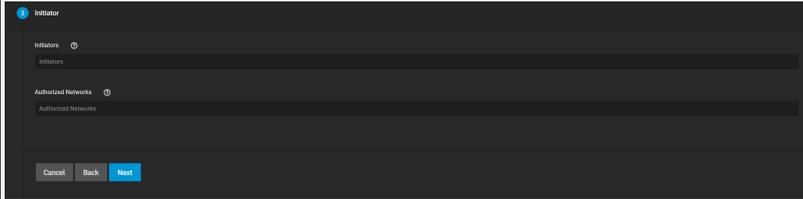
Initiator

100%

Text Edge Style

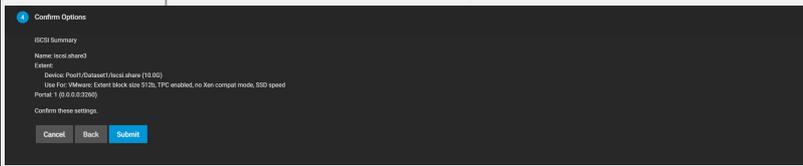
None

Decide which initiators or networks can use the iSCSI share. Leave the list empty to allow all initiators or networks, or add entries to the list to limit access to those systems.



Confirm

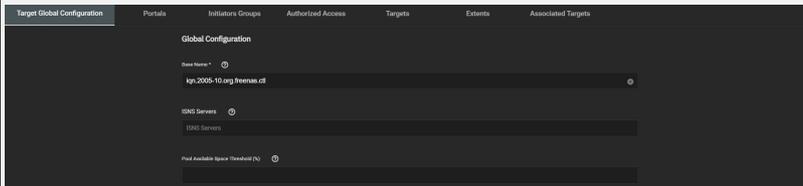
Confirm the settings are correct and click **Submit**.



Manual Setup

Target Global Configuration

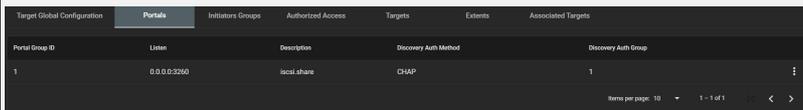
The **Target Global Configuration** tab lets users configure settings that apply to all iSCSI shares.



Setting	Description
Base Name	Lowercase alphanumeric characters plus dot (.), dash (-), and colon (:) are allowed. See the "Constructing iSCSI names using the iqn.format" section of RFC3721 .
ISNS Servers	Hostnames or IP addresses of the ISNS servers to register with the iSCSI targets and portals of the system. Separate entries by pressing Enter.
Pool Available Space Threshold (%)	Generates an alert when the pool has this percent space remaining. This is typically configured at the pool level when using zvols or at the extent level for both file and device-based extents.

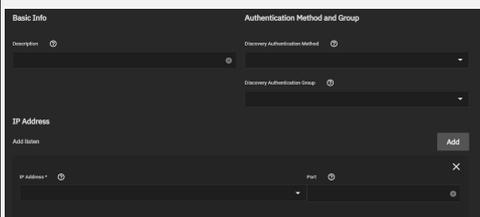
Portals

The **Portals** tab lets users create new portals or edit existing ones in the list.



To add a new portal, click **ADD** and enter the basic and IP address information.

To edit an existing portal, click  next to the portal and select **Edit**.



Basic Info

Setting	Description
Description	Optional description. Portals are automatically assigned a numeric group.

Authentication Method and Group

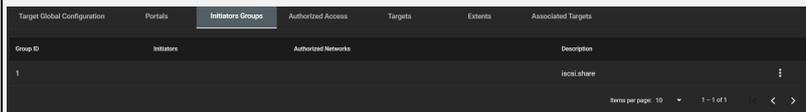
Setting	Description
Discovery Authentication Method	iSCSI supports multiple authentication methods that the target uses to discover valid devices. None allows anonymous discovery while CHAP and Mutual CHAP require authentication.
Discovery Authentication Group	Group ID created in Authorized Access . Required when the Discovery Authentication Method is CHAP or Mutual CHAP.

IP Address

Setting	Description
IP Address	Select the IP addresses the portal listens to. Click Add to add IP addresses with a different network port. 0.0.0.0 listens on all IPv4 addresses and :: listens on all IPv6 addresses.
Port	TCP port used to access the iSCSI target. Default is 3260 .
ADD	Adds another IP address row.

Initiators Groups

The **Initiators Groups** tab lets users create new authorized access client groups or edit existing ones in the list.



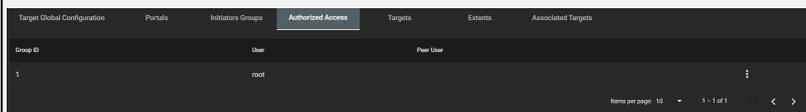
To add a new initiators group, click **Add** and leave **Allow All Initiators** checked or configure your own allowed initiators and authorized networks.

To edit an existing initiators group, click  next to the initiators group and select **Edit**.

Setting	Description
Allow All Initiators	Allows all initiators when checked.
Allowed Initiators (IQN)	Initiators allowed access to this system. Enter an iSCSI Qualified Name (IQN) and click + to add it to the list. Example: <i>iqn.1994-09.org.freebsd:freenas.local</i> .
Authorized Networks	Network addresses allowed to use this initiator. Each address can include an optional CIDR netmask. Click + to add the network address to the list. Example: <i>192.168.2.0/24</i> .
Description	Any notes about initiators.

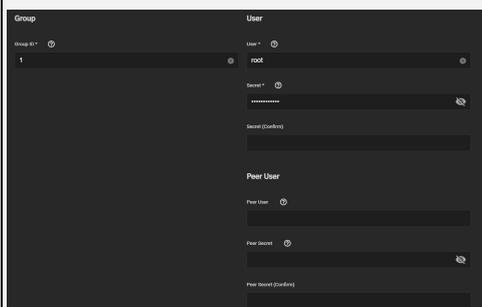
Authorized Access

The **Authorized Access** tab lets users create new authorized access networks or edit existing ones in the list.



To add a new authorized access network, click **ADD** and fill out the group, user, and peer user information.

To edit an existing authorized access network, click  next to it and select **Edit**.



Group

Setting	Description
Group ID	Allow configuring different groups with different authentication profiles. Example: all users with a group ID of <i>1</i> inherits the authentication profile associated with <i>Group 1</i> .

User

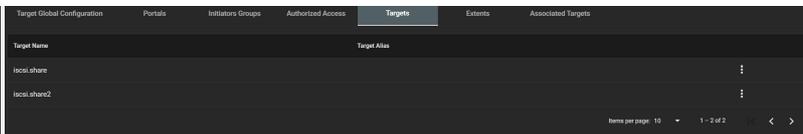
Setting	Description
User	User account to create for CHAP authentication with the user on the remote system. Many initiators use the initiator name as the user name.
Secret	User password. Must be at least 12 and no more than 16 characters long.
Secret (Confirm)	Confirm the user password.

Peer User

Setting	Description
Peer User	Only entered when configuring mutual CHAP. Usually the same value as User .
Peer Secret	Mutual secret password. Required when Peer User is set. Must be different than the Secret .
Peer Secret (Confirm)	Confirm the mutual secret password.

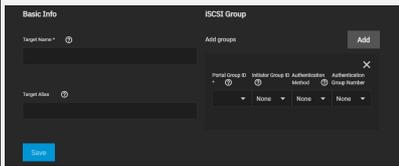
Targets

The **Targets** tab lets users create new TrueNAS storage resources or edit existing ones in the list.



To add a new target, click **ADD** and enter the basic and iSCSI group information.

To edit an existing target, click **Edit** next to it and select **Edit**.



Basic Info

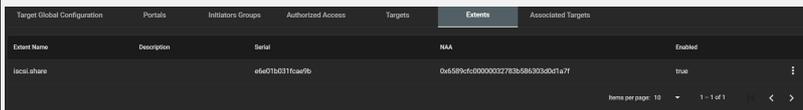
Setting	Description
Target Name	The base name is automatically prepended if the target name does not start with iqn. Lowercase alphanumeric characters plus dot (.), dash (-), and colon (:) are allowed. See the "Constructing iSCSI names using the iqn.format" section of RFC3721 .
Target Alias	Optional user-friendly name.

iSCSI Group

Setting	Description
Portal Group ID	Leave empty or select an existing portal to use.
Initiator Group ID	Select the existing initiator group that has access to the target.
Authentication Method	Choices are None , Auto , CHAP , or Mutual CHAP .
Authentication Group Number	Select None or an integer. This value represents the number of existing authorized accesses.

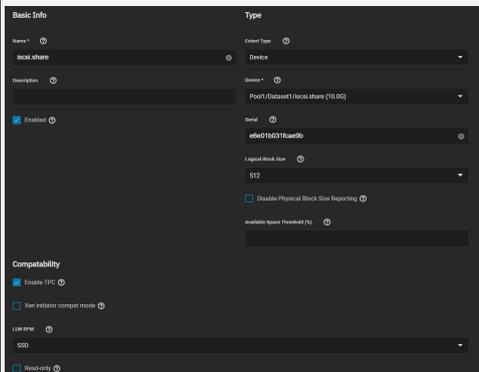
Extents

The **Extents** tab lets users create new shared storage units or edit existing ones in the list.



To add a new extent, click **Add** and enter the information.

To edit an existing extent, click **Edit** next to it and select **Edit**.



Basic Info

Setting	Description
Name	Name of the extent. If the Extent size is not 0 , it cannot be an existing file within the pool or dataset.
Description	Notes about this extent.
Enabled	Set to enable the iSCSI extent.

Type

Setting	Description
Extent Type	Device provides virtual storage access to zvols, zvol snapshots, or physical devices. File provides virtual storage access to a single file.
Device	Only appears if Device is selected. Select the unformatted disk, controller, or zvol snapshot.
Path to the Extent	Only appears if File is selected. Browse to an existing file. Create a new file by browsing to a dataset and appending <code>/{filename.ext}</code> to the path. Users cannot create extents inside a jail root directory.
Filesize	Only appears if File is selected. Entering 0 uses the actual file size and requires that the file already exists. Otherwise, specify the file size for the new file.
Logical Block Size	Leave at the default of 512 unless the initiator requires a different block size.
Disable Physical Block Size Reporting	Set if the initiator does not support physical block size values over 4K (MS SQL).

Compatibility

Setting	Description
Enable TPC	Set to allow an initiator to bypass normal access control and access any scannable target. This allows xcopy operations that are otherwise blocked by access control.
Xen initiator compat mode	Set when using Xen as the iSCSI initiator.
LUN RPM	Do <i>NOT</i> change this setting when using Windows as the initiator. You only change it in large environments where the number of systems

	using a specific RPM is needed for accurate reporting statistics.
Read-only	Set to prevent the initiator from initializing this LUN.

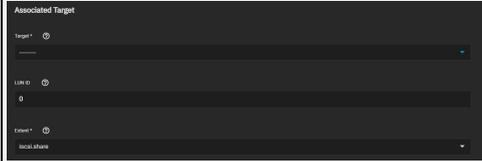
Associated Targets

The **Associated Targets** tab lets users create new associated TrueNAS storage resources or edit existing ones in the list.



To add a new associated target, click **ADD** and fill out the information.

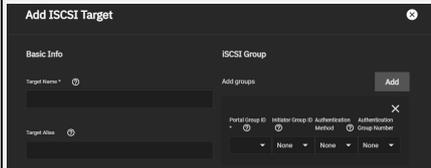
To edit an existing associated target, click  next to it and select **Edit**.



Setting	Description
Target	Select an existing target.
LUN ID	Select the value or enter a value between 0 and 1023. Some initiators expect a value below 256. Leave this field blank to automatically assign the next available ID.
Extent	Select an existing extent.

Quick iSCSI Target Creation

TrueNAS SCALE allows users to add iSCSI targets without having to set up another share. Go to **Shares** and click **Add** in the **Block (iSCSI) Shares Targets** window.



Basic Info

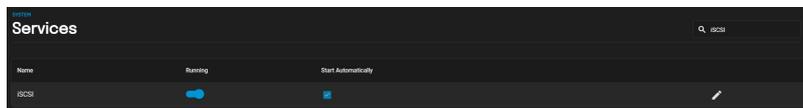
Setting	Description
Target Name	The base name is automatically prepended if the target name does not start with iqn. Lowercase alphanumeric characters plus dot (.), dash (-), and colon (:) are allowed. See the "Constructing iSCSI names using the iqn.format" section of RFC3721 .
Target Alias	Optional user-friendly name.

iSCSI Group

Setting	Description
Portal Group ID	Leave empty or select an existing portal to use.
Initiator Group ID	Select which existing initiator group has access to the target.
Authentication Method	Choices are None , Auto , CHAP , or Mutual CHAP .
Authentication Group Number	Select None or an integer. This value represents the number of existing authorized accesses.

Starting the iSCSI Service

To turn on the iSCSI service, go to **Services** and toggle on **iSCSI**. Set **Start Automatically** to start it when TrueNAS boots up.



Clicking the  returns to the options in **Shares > Block (iSCSI) Shares Targets**.

Using the iSCSI Share

Connecting to and using an iSCSI share can differ between operating systems.

Linux

iSCSI Utilities and Service

First, open the command line and ensure you have installed the `open-iscsi` utility. To install the utility on an Ubuntu/Debian distribution, enter `sudo apt update && sudo apt install open-iscsi`. After the installation completes, ensure the `iscsid` service is running: `sudo service iscsid start`. With the `iscsid` service started, run the `iscsiadm` command with the discovery arguments and get the necessary information to connect to the share.

```

cruenas@LinuxMachine:~$ sudo apt update && sudo apt install open-iscsi
Hit:1 http://us.archive.ubuntu.com/ubuntu focal InRelease
Hit:1 http://us.archive.ubuntu.com/ubuntu focal-updates InRelease
Hit:1 http://us.archive.ubuntu.com/ubuntu focal-backports InRelease
Hit:1 http://us.archive.ubuntu.com/ubuntu focal-security InRelease
Reading package lists... Done

```

Discover and Log In to the iSCSI Share

Run the command `sudo iscsiadm \--mode discovery \--type sendtargets \--portal {IPADDRESS}`. The output provides the basename and target name that TrueNAS configured.

```
truenas@LinuxMachine:~$ sudo iscsiadm \--mode discovery \--type sendtargets \--portal 10.10.10.10 238.15.118.3260,-1 iqn.2005-10.org.freenas.ctl:iscsishare
10.238.15.118.3260,-1 iqn.2005-10.org.freenas.ctl:iscsishare2
10.238.15.118.3260,-1 iqn.2005-10.org.freenas.ctl:iscsifile
truenas@LinuxMachine:~$
```

Alternatively, enter `sudo iscsiadm -m discovery -t st -p {IPADDRESS}` to get the same output. Note the basename and target name given in the output. You need them to log in to the iSCSI share.

When a Portal Discovery Authentication Method is CHAP, add the three following lines to `/etc/iscsi/iscsid.conf`.

```
discovery.sendtargets.auth.authmethod = CHAP
discovery.sendtargets.auth.username = user
discovery.sendtargets.auth.password = secret
```

The user for `discovery.sendtargets.auth.username` is set in the **Authorized Access** used by the iSCSI share **Portal**. Likewise, the password to use for `discovery.sendtargets.auth.password` is the **Authorized Access** secret. Without those lines, the `iscsiadm` does not discover the portal with the CHAP authentication method.

Next, enter `sudo iscsiadm \--mode node \--targetname {BASENAME}:{TARGETNAME} \--portal {IPADDRESS} \--login`, where `{BASENAME}` and `{TARGETNAME}` is the `discovery` command information.

```
truenas@LinuxMachine:~$ sudo iscsiadm \--mode discovery \--type sendtargets \--portal freenas.local
freenas.local:3260,-1 iqn.2005-10.org.freenas.ctl:iscsi.share
truenas@LinuxMachine:~$ sudo iscsiadm \--mode node \--targetname iqn.2005-10.org.freenas.ctl:iscsi.share \--portal freenas.local \--login
Login in to [iface: default, target: iqn.2005-10.org.freenas.ctl:iscsi.share, portal: freenas.local:3260] (multiple)
Login to [iface: default, target: iqn.2005-10.org.freenas.ctl:iscsi.share, portal: freenas.local:3260] successful.
truenas@LinuxMachine:~$
```

Partition iSCSI Disk

When the iSCSI share login succeeds, the device shared through iSCSI shows on the Linux system as an **iSCSI Disk**. To view a list of connected disks in Linux, enter `sudo fdisk -l`.

```
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk /dev/sda: 476.96 GiB, 512110190592 bytes, 1000215216 sectors
Disk model: SAMSUNG MZNLN512
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: gpt
Disk identifier: B7D9E3B0-EBED-4CEA-9CC6-08F2918A54FB

Device      Start      End          Sectors    Size Type
/dev/sda1   2048       1050623     1048576    512M EFI System
/dev/sda2  1050624   1000214527 999163904 476.4G Linux filesystem

Disk /dev/loop8: 240.82 MiB, 252493824 bytes, 493152 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk /dev/loop9: 29.84 MiB, 31272960 bytes, 61080 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk /dev/sdb: 10 GiB, 10737434624 bytes, 2621444 sectors
Disk model: iSCSI Disk
Units: sectors of 1 * 4096 = 4096 bytes
Sector size (logical/physical): 4096 bytes / 16384 bytes
I/O size (minimum/optimal): 16384 bytes / 1048576 bytes
truenas@LinuxMachine:~$
```

Because the connected iSCSI disk is raw, you must partition it. Identify the iSCSI device in the list and enter `sudo fdisk {/PATH/TO/ISCSIDEVICE}`.

```
truenas@LinuxMachine:~$ sudo fdisk /dev/sdb

Welcome to fdisk (util-linux 2.34).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): n
Partition type
  p   primary (0 primary, 0 extended, 4 free)
  e   extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1):
First sector (256-2621443, default 256):
Last sector, +/-sectors or +/-size[K,M,G,T,P] (256-2621443, default 2621443):

Created a new partition 1 of type 'Linux' and of size 10 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.

truenas@LinuxMachine:~$
```

Shell lists the iSCSI device path in the `sudo fdisk -l` output. Use the `fdisk` command defaults when partitioning the disk.

Remember to type `w` when finished partitioning the disk. The `w` command tells `fdisk` to save any changes before quitting.

```
truenas@LinuxMachine:~$ sudo mkfs /dev/sdb1
mke2fs 1.45.5 (07-Jan-2020)
Discarding device blocks: done
Creating filesystem with 2621188 4k blocks and 655360 inodes
Filesystem UUID: 1b38f07a-bb23-40ab-b1eb-255480e4dbbc
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Writing superblocks and filesystem accounting information: done

truenas@LinuxMachine:~$
```

After creating the partition on the iSCSI disk, a partition slice displays on the device name. For example, `/dev/sdb1`. Enter `fdisk -l` to see the new partition slice.

Make a Filesystem on the iSCSI Disk

Finally, use `mkfs` to make a filesystem on the device's new partition slice. To create the default filesystem (ext2), enter `sudo mkfs {/PATH/TO/ISCSIDEVICEPARTITIONSLICE}`.

```
truenas@LinuxMachine:~$ sudo mkfs /dev/sdb1
mke2fs 1.45.5 (07-Jan-2020)
Discarding device blocks: done
Creating filesystem with 2621188 4k blocks and 655360 inodes
Filesystem UUID: 1b30f07a-bb23-40ab-b1eb-255480e4dbbc
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Writing superblocks and filesystem accounting information: done
```

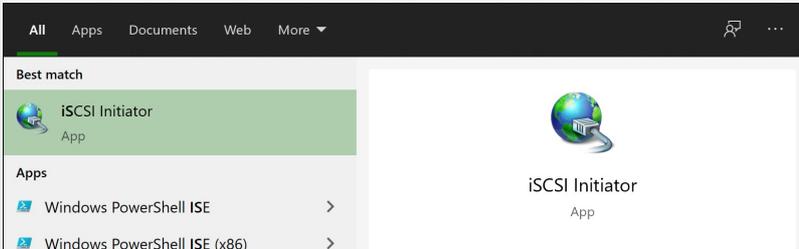
Mount the iSCSI Device

Now the iSCSI device can mount and share data. Enter `sudo mount {/PATH/TO/iSCSIDevicePARTITIONSLICE}`. For example, `sudo mount /dev/sdb1 /mnt` mounts the iSCSI device `sdb1` to `/mnt`.

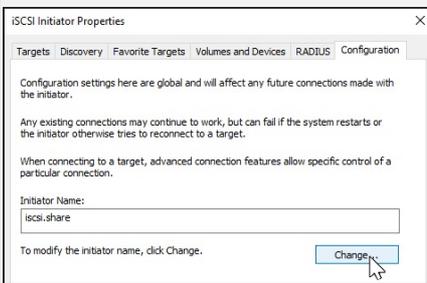
Windows

To access the data on the iSCSI share, clients need to use iSCSI Initiator software. An iSCSI Initiator client is pre-installed in Windows 7 to 10 Pro, and Windows Server 2008, 2012, and 2019. Windows Professional Edition is usually required.

First, click the **Start Menu** and search for the **iSCSI Initiator** application.



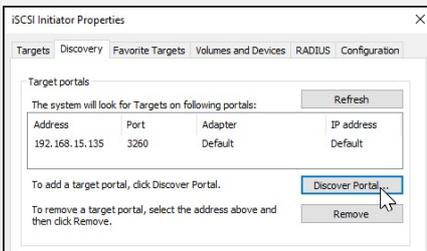
Next, go to the **Configuration** tab and click **Change** to replace the iSCSI initiator with the name created earlier. Click **OK**.



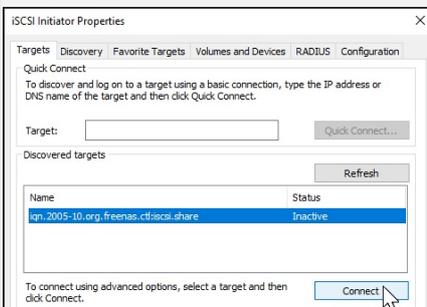
Next, switch to the **Discovery** Tab, click **Discover Portal**, and type in the TrueNAS IP address.

- If TrueNAS changed the port number from the default **3260**, enter the new port number.
- If you set up CHAP when creating the iSCSI share, click **Advanced...**, set **Enable CHAP log on**, and enter the initiator name and the same target/secret set earlier in TrueNAS.

Click **OK**.

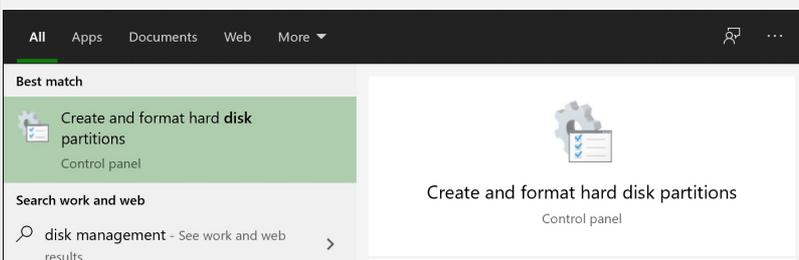


Go to the **Targets** tab, highlight the iSCSI target, and click **Connect**.

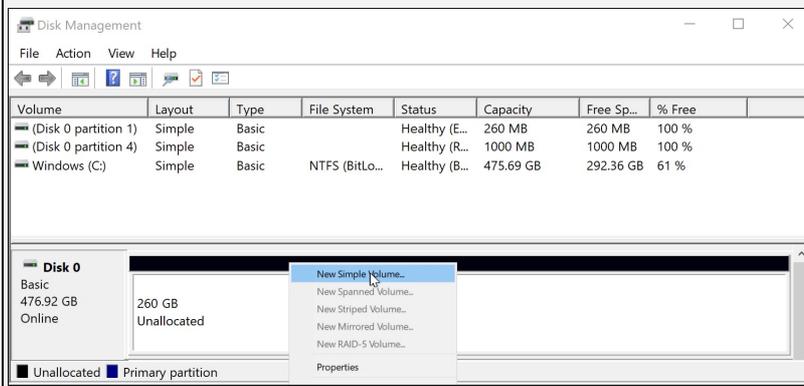


After Windows connects to the iSCSI target, you can partition the drive.

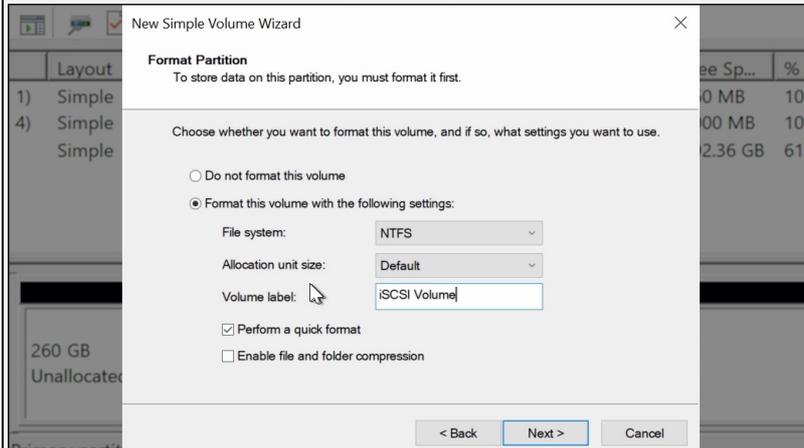
Search for and open the **Disk Management** app.



The current state of your drive should be **unallocated**. Right-click the drive and click **New Simple Volume...**



Complete the Wizard to format the drive and assign a drive letter and name.



Finally, go to **This PC** or **My Computer** in **File Explorer**. The new iSCSI volume should display under the list of drives. You should now be able to add, delete, and modify files and folders on your iSCSI drive.

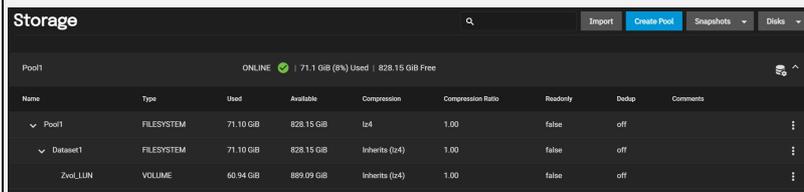


Expanding LUNs

TrueNAS lets users expand Zvol and file-based LUNs to increase the available storage that the iSCSI shares.

Zvol LUN

To expand a Zvol LUN, go to **Storage** and click the **Edit Zvol** next to the Zvol LUN, then select **Edit Zvol**.



Enter a new size in the **Size for this zvol** field, then click **SAVE**.

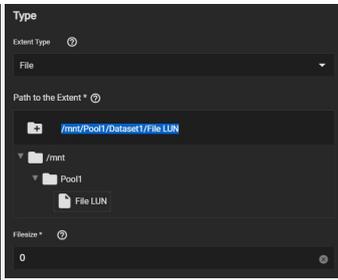


TrueNAS prevents data loss by not allowing users to reduce the Zvol size. TrueNAS also does not allow users to increase the Zvol size past 80% of the pool size.

File LUN

You need to know the path to the file to expand a file-based LUN. Go to **Shares** and click **Configure** in the **Block (iSCSI) Shares Targets** window, then select the **Extents** tab.

Click the **Edit** next to the file-based LUN and select **Edit**.



Highlight and copy the path, then click **Cancel**.

Go to **Shell** and input `truncate -s +[size] [path to file]`, then press Enter.

The *[size]* is how much space you want to grow the file by, and the *[path to file]* is the file path you copied earlier.

```
Linux truenas.scale.local 5.10.42+truenas #1 SMP Mon Aug 30 21:54:59 UTC 2021 x86_64

TrueNAS (c) 2009-2021, iXsystems, Inc.
All rights reserved.
TrueNAS code is released under the modified BSD license with some files copyrighted by (c) iXsystems, Inc.

For more information, documentation, help or support, go here: http://truenas.com
Welcome to TrueNAS
Last login: Fri Sep 17 06:52:53 PDT 2021 on pts/2
root@truenas# truncate -s +2g /mnt/Pool1/Dataset1/File LUN
```

An example command could look like this: `truncate -s +2g /mnt/Pool1/Dataset1/File LUN`

Lastly, go back to the extent in **Shares > Block (iSCSI) Shares Targets** and make sure the **Filesize** is set to **0** so that the share uses the actual file size.

Windows (SMB) Shares

SMB Shares



Video Player is loading.
<https://www.truenas.com/docs/files/scaleangelfishsmbshare.mp4>

Play Video

Play

SMB (also known as CIFS) is the native file sharing system in Windows. SMB shares can connect to most operating systems, including Windows, macOS, and Linux. TrueNAS can use SMB to share files among single or multiple users or devices.

SMB supports a wide range of permissions, security settings, and advanced permissions (ACLs) on Windows and other systems, as well as Windows Alternate Streams and Extended Metadata. SMB is suitable for managing and administering large or small pools of data.

TrueNAS uses Samba to provide SMB services. The SMB protocol has multiple versions. An SMB client typically negotiates the highest supported SMB protocol during SMB session negotiation. In the industry-wide, SMB1 protocol (sometimes referred to as NT1) usage is **being deprecated** for security reasons. However, most SMB clients support SMB 2 or 3 protocols, even when they are not default.

Remaining Time -1:31

Legacy SMB clients rely on NetBIOS name resolution to discover SMB servers on a network. TrueNAS disables the NetBIOS Name Server (nmbd) by default. Enabled in **Network** if you require its functionality.

macOS clients use **mDNS** to discover SMB servers present on the network. TrueNAS enables the mDNS server (avahi) by default.

Windows clients use **WS-Discovery** to discover the presence of SMB servers, but network discovery can be disabled by default depending on the Windows client version.

Discoverability through broadcast protocols is a convenience feature and not required to access an SMB server.

First Steps

- descriptions off, selected

Create a Dataset

It is recommended to create a new dataset and set the **Share Type** to **SMB** for the new SMB share.

- captions off, selected

What does this do?

TrueNAS creates the ZFS dataset with these settings:

- **aclmode = restricted**
- **case sensitivity = insensitive**

TrueNAS also applies a default access control list to the dataset. This default ACL is restrictive and only allows access to the dataset owner and group. You can modify the ACL in a console window.

Beginning of dialog window. Escape will cancel and close the window.

Create Local User Accounts

By default, all new local users are members of a built-in SMB group called **builtin users**. You can use the group to grant access to all local users on the server or add more groups to fine-tune permissions to large numbers of users. You cannot access SMB shares with user accounts built-in to TrueNAS or those without the **smb** flag.

Why not just allow anonymous access to the share?

Anonymous or guest access to the share is possible, but it is a security vulnerability. Major SMB client vendors are deprecating it, partly because signing and encryption are not possible for guest access.

What about LDAP users?

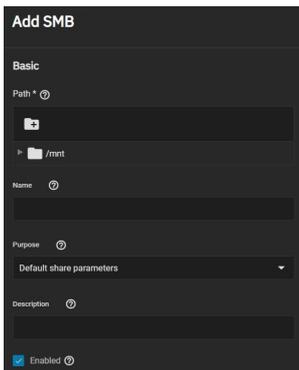
If you want LDAP server users to access the SMB share, go to **Credentials > Directory Services**. Click **Settings** in the **LDAP window**, then click **Advanced Options and Samba Schema**. Local TrueNAS user accounts no longer have access to the share.

Tune the Dataset ACL

After creating a dataset and accounts, you need to investigate your access requirements and adjust the dataset ACL to match. Go to **Storage**, open the options for the new dataset, and click **Edit Permissions**. Many home users typically add a new entry that grants **FULL CONTROL** to the **builtin_users** group with the flags set to **INHERIT**. See the [Permissions article](#) for more details.

Creating the SMB Share

To create a Windows SMB share, go to **Sharing > Windows Shares (SMB)** and click **Add**.



The SMB share **Path** and **Name** define the absolute minimum amount of information required to create a new SMB share. The **Path** is the directory tree on the local filesystem that TrueNAS exports over the SMB protocol. The **Name** is the SMB share name, which forms part of the full share pathname when SMB clients perform an SMB tree connect. Because of how the SMB protocol uses the name, it must be less than or equal to 80 characters. It cannot have any invalid characters as specified in Microsoft documentation MS-FSCC section 2.1.6. If a name is not supplied, then the last component of the path is used as the share name.

You can set a share **Purpose** to apply and lock pre-determined advanced options for the share. To retain control over all the share **Advanced Options**, choose **No presets**.

What do all the presets do?

The following table shows the preset options for the different **Purpose** options and if those are locked.

A [] indicates the option is enabled while [] indicates the option is disabled. [] indicates empty text fields, and [%U] indicates the exact option the preset created. [] means the option is disabled.

Setting	Default share parameters	Multi-user time machine	Multi-pr
Enable ACL	(locked)		(locked)
Export Read Only	(locked)		
Browsable to Network Clients	(locked)		
Allow Guest Access			
Access Based Share Enumeration	(locked)		
Hosts Allow	(locked)		
Hosts Deny	(locked)		
Use as Home Share	(locked)		

Time Machine	(locked)		
Enable Shadow Copies	(locked)		
Export Recycle Bin	(locked)		
Use Apple-style Character Encoding	(locked)		
Enable Alternate Data Streams	(locked)		(locked)
Enable SMB2/3 Durable Handles	(locked)		(locked)
Enable FSRVP	(locked)		
Path Suffix	[] (locked)	[%U] (locked)	[%U]
Auxiliary Parameters	[]	[]	[]

You can specify an optional **Description** to help explain the share's purpose.

Enabled allows this path to be shared when the SMB service is activated. Unsetting **Enabled** disables the share without deleting the configuration.

Advanced Options

Access

- Enable ACL
- Export Read Only
- Browsable to Network Clients
- Allow Guest Access
- Access Based Share Enumeration

Hosts Allow

Hosts Allow

Hosts Deny

Hosts Deny

Other Options

- Use as Home Share
- Time Machine
- Legacy AFP Compatibility
- Enable Shadow Copies
- Export Recycle Bin
- Use Apple-style Character Encoding
- Enable Alternate Data Streams
- Enable SMB2/3 Durable Handles
- Enable FSRVP

Path Suffix

Auxiliary Parameters

Options are divided into **Access** and **Other Options** groups. **Access** options control various settings for allowing systems or users to access or modify the shared data.

Setting	Description
Enable ACL	Enables ACL support for the SMB share.
Export Read Only	Prohibits writes to the share.
Browsable to Network Clients	Determine whether this share name is included when browsing shares. Home shares are only visible to the owner regardless of this setting.
Allow Guest Access	Privileges are the same as the guest account. Guest access is disabled by default in Windows 10 version 1709 and Windows Server version 1903. Additional client-side configuration is required to provide guest access to these clients. MacOS clients: Attempting to connect as a user that does not exist in FreeNAS <i>does not</i> automatically connect as the guest account. The Connect As: Guest option must be specifically chosen in macOS to log in as the guest account. See the Apple documentation for more details.
Access Based Share Enumeration	Restrict share visibility to users with read or write access to the share. See the smb.conf manual page.
Hosts Allow	Enter a list of allowed hostnames or IP addresses. Separate entries by pressing Enter. You can find a more detailed description with examples here .
Hosts Deny	Enter a list of denied hostnames or IP addresses. Separate entries by pressing Enter.

The **Hosts Allow** and **Hosts Deny** fields work together to produce different situations:

- If neither **Hosts Allow** or **Hosts Deny** contains an entry, then SMB share access is allowed for any host.
- If there is a Hosts Allow list but no Hosts Deny list, then only allow hosts on the Hosts Allow list.
- If there is a Hosts Deny* list but no Hosts Allow list, then allow all hosts on the Hosts Deny list.
- If there is both a Hosts Allow and Hosts Deny list, then allow all hosts on the Hosts Allow list. If there is a host not on the Hosts Allow and not on the Hosts Deny list, then allow it.

The **Other Options** have settings for improving Apple software compatibility, ZFS snapshot features, and other advanced features.

Setting	Description
Use as Home Share	Allows the share to host user home directories. Each user is given a personal home directory when connecting to the share which is not accessible by other users. This allows for a personal, dynamic share. Only one share can be used as the home share. See the SMB Home Shares section below.
Time Machine	Enables Apple Time Machine backups on this share.
Legacy AFP Compatibility	This controls how the SMB share reads and writes data. Leave unset for the share to behave like a normal SMB share and set for the share to behave like the deprecated Apple Filing Protocol (AFP). Only set this when this share originated as an AFP sharing configuration. This is not required for pure SMB shares or macOS SMB clients.
Enable Shadow Copies	Export ZFS snapshots as Shadow Copies for Microsoft Volume Shadow Copy Service (VSS) clients.
Export Recycle Bin	Files that are deleted from the same dataset are moved to the Recycle Bin and do not take any additional space. <i>Deleting files over NFS removes the files permanently.</i> When the files are in a different dataset or a child dataset, they are copied to the dataset where the recycle bin is located. To prevent excessive space usage, files larger than 20 MiB are deleted rather than moved. Adjust the Auxiliary Parameter <code>crossrename: sizeLimit=</code> setting to allow larger files. For example, <code>crossrename: sizeLimit=50</code> allows moves of files up to 50 MiB in size. This means

	files can be permanently deleted or moved from the recycle bin. <i>This is not a replacement for ZFS snapshots.</i>
Use Apple-style Character Encoding	By default, Samba uses a hashing algorithm for NTFS illegal characters. Enabling this option converts NTFS illegal characters in the same manner as macOS SMB clients.
Enable Alternate Data Streams	Allows multiple NTFS data streams . Disabling this option causes macOS to write streams to files on the filesystem.
Enable SMB2/3 Durable Handles	Allow using open file handles that can withstand short disconnections. Support for POSIX byte-range locks in Samba is also disabled. This option is not recommended when configuring multi-protocol or local access to files.
Enable FSRVP	Enable support for the File Server Remote VSS Protocol (FSVRP). This protocol allows remote procedure call (RPC) clients to manage snapshots for a specific SMB share. The share path must be a dataset mount point. Snapshots have the prefix <code>fs-</code> followed by a snapshot creation timestamp. A snapshot must have this prefix for an RPC user to delete it.
Path Suffix	Appends a suffix to the share connection path. This is used to provide unique shares on a per-user, per-computer, or per-IP address basis. Suffixes can contain a macro. See the smb.conf manual page for a list of supported macros. The connectpath must be preset before a client connects.
Auxiliary Parameters	Additional smb.conf settings.

Clicking **Save** creates the share and adds it to the **Shares > Windows (SMB) Shares** list. You can also choose to enable the SMB service at this time.

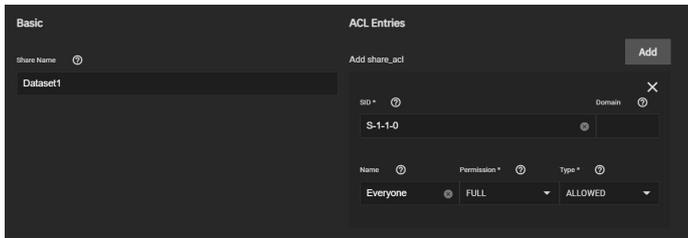
Share Management

After creating the SMB share, additional management options are available by going to the **Shares** screen and clicking **Windows (SMB) Shares** in the **Windows (SMB) Shares** window. Click **⋮** next to the share you want to manage.

- **Edit:** Opens the [share creation screen](#) to reconfigure the share or disable it.
- **Edit Share ACL:** Opens a screen to configure an Access Control List (ACL) for the share. This screen is separate from filesystem permissions and applies at the entire SMB share level. Other filesharing protocol clients or other SMB shares that export the same share **Path** do not interpret these permissions. Open is the default. This ACL determines the browse list if **Access Based Share Enumeration** is enabled.
- **Edit Filesystem ACL:** Opens a screen to configure an Access Control List (ACL) for the path defined in the share **Path**.
- **Delete:** Remove the share configuration from TrueNAS. Shared data is unaffected.

Configure Share ACL

To see the share ACL options, select **Edit Share ACL**.



TrueNAS displays the **Share Name** (cannot be changed). **ACL Entries** are listed as a block of settings. Click **Add** to register a new entry.

Setting	Description
SID	Who this ACL entry (ACE) applies to, shown as a Windows Security Identifier . Values in either SID or Domain with Name is required for the ACL.
Domain	Domain for the user specified in Name . Required when a SID value is not entered. Local users have the SMB server NetBIOS name: <code>truenas\smusers</code> .
Permission	Predefined permission combinations. Read: Read access and Execute permission on the object (RX). Change: Read access, Execute permission, Write access, and Delete object (RXWD). Full: Read access, Execute permission, Write access, Delete object, change Permissions, and take Ownership (RXWDPO). For more details, see smbacls(1) .
Name	Who this ACL entry applies to, shown as a user name. Requires adding the user Domain .
Type	How permissions are applied to the share. Allowed denies all permissions by default except those that are manually defined. Denied allows all permissions by default except those that are manually defined.

Clicking **Save** stores the share ACL and applies it to the share immediately.

Configure Filesystem ACL

Selecting **Edit Filesystem ACL** takes you to the **Edit File ACL** screen in **Storage** to edit the dataset ACL.

Since SCALE gives users the option to use either POSIX or NFSv4 share [ACL types](#), the **Edit File ACL** page differs depending on which ACL type the filesystem is using.

NFSv4 Filesystem ACL

The filesystem ACL defines the user accounts or groups that own or have specific [permissions](#) to the shared dataset. The **User** and **Group** values show which accounts *own* or have full permissions to the dataset. Change the default settings to your preferred primary account and group and set the **Apply permissions recursively** checkbox before saving any changes.

ACL Presets

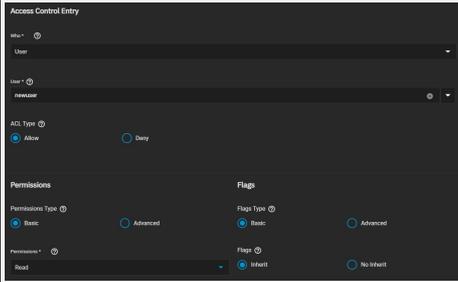
To rewrite the current ACL with a standardized preset, click **Use ACL Preset** and choose an option:

NFS4_OPEN: Owner and group have full dataset control. All other accounts can modify the dataset contents.

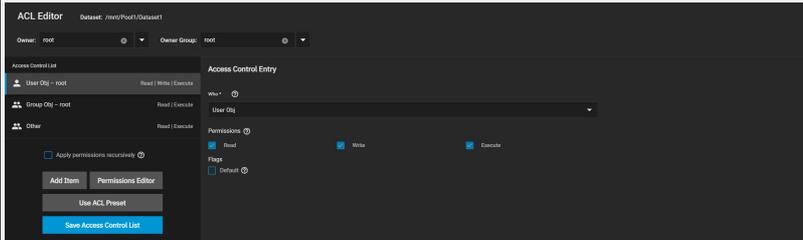
NFS4_RESTRICTED: Owner has full dataset control. Group can modify the dataset contents. **NFS4_HOME:** Owner has full dataset control. Group can modify the dataset contents. All other accounts can navigate the dataset.

Adding ACL Entries (ACEs)

To define permissions for a specific user account or group, click **Add Item**. Open the **Who** drop-down, select **User** or **Group**, and choose a specific user or group account. Define how the settings apply to the account, then choose which permissions to apply. For example, to only allow the *newuser* user permission to view dataset contents but not make changes, set the **ACL Type** to **Allow** and **Permissions** to **Read**.



POSIX Filesystem ACL



The filesystem ACL defines the user accounts or groups that own or have specific [permissions](#) to the shared dataset.

The **User** and **Group** values show which accounts own, or have full permissions to the dataset. Change the default settings to your preferred primary account and group and set the **Apply permissions recursively** checkbox before saving any changes.

ACL Presets

To rewrite the current ACL with a standardized preset, click **Use ACL Preset** and choose an option:

POSIX_OPEN: Owner and group have full dataset control. All other accounts can modify the dataset contents.

POSIX_RESTRICTED: Owner has full dataset control. Group can modify the dataset contents. **POSIX_HOME:** Owner has full dataset control. Group can modify the dataset contents. All other accounts can navigate the dataset.

Adding ACL Entries (ACEs)

To define permissions for a specific user account or group, click **Add Item**. Open the **Who** drop-down, select **User** or **Group**, and choose a specific user or group account. Define how the settings apply to the account, then choose which permissions to apply. For example, to only allow the *newuser* user permission to view dataset contents but not make changes, set the **ACL Type** to **Allow** and **Permissions** to **Read**.



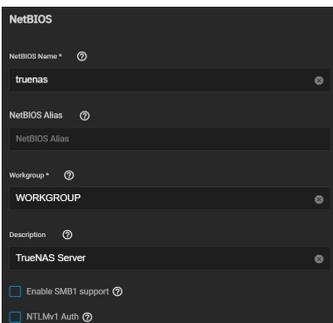
Activate the SMB Service

Connecting to an SMB share does not work when the related system service is not activated. To make SMB share available on the network, go to **System Settings > Services** and click the toggle to running for **SMB**. Set **Start Automatically** if you want the service to activate when TrueNAS boots.

Service Configuration

Configure the SMB service by clicking

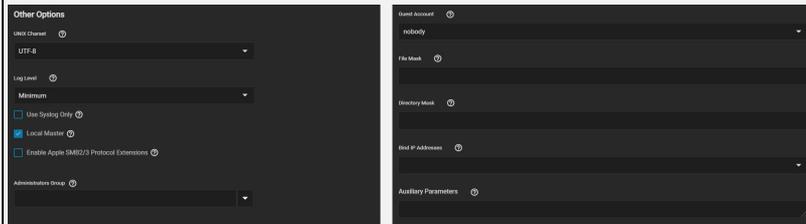
. Unless you need a specific setting or are configuring a unique network environment, we recommend the default settings.



Setting	Description
NetBIOS Name	Automatically populated with the system's original hostname. This name is limited to 15 characters and cannot be the Workgroup name.
NetBIOS Alias	Enter any aliases, separated by spaces. Each alias can be up to 15 characters long.
Workgroup	Must match the Windows workgroup name. When this is unconfigured and Active Directory or LDAP is active, TrueNAS detects and sets the correct workgroup from these services.
Description	This allows entering any notes or descriptive details about the service configuration.
Enable SMB1	Allow legacy SMB1 clients to connect to the server. Note that SMB1 is being deprecated and it is advised to upgrade clients to operating system versions that support modern SMB protocol versions.

support	
NTLMv1 Auth	When set, smbd attempts to authenticate users with the insecure and vulnerable NTLMv1 encryption. This setting allows backward compatibility with older versions of Windows, but is not recommended and should not be used on untrusted networks.

Advanced Options



Setting	Description
UNIX Charset	Character set used internally. UTF-8 is standard for most systems as it supports all characters in all languages.
Log Level	Record SMB service messages up to the specified log level. By default, error and warning level messages are logged. It is not recommended to use a log level above minimum for production servers.
Use Syslog Only	Set to log authentication failures in <code>/var/log/messages</code> instead of the default <code>/var/log/samba4/log.smbd</code> .
Local Master	Set to determine if the system participates in a browser election. Clear the setting when the network contains an AD or LDAP server, or when Vista or Windows 7 machines are present.
Enable Apple SMB2/3 Protocol Extensions	These protocol extensions can be used by macOS to improve the performance and behavioral characteristics of SMB shares. This is required for Time Machine support.
Administrators Group	Members of this group are local administrators and automatically have privileges to take ownership of any file in an SMB share, reset permissions, and administer the SMB server through the Computer Management MMC snap-in.
Guest Account	Account used for guest access. Default is nobody . The chosen account is required to have permissions to the shared pool or dataset. To adjust permissions, edit the dataset Access Control List (ACL), add a new entry for the chosen guest account, and configure the permissions in that entry. If the selected Guest Account is deleted the field resets to nobody .
File Mask	Overrides default 0666 file creation mask which creates files with read and write access for everybody.
Directory Mask	Overrides default directory creation mask of 0777 which grants directory read, write and execute access for everybody.
Bind IP Addresses	Static IP addresses which SMB listens on for connections. Leaving all unselected defaults to listening on all active interfaces.
Auxiliary Parameters	Stores additional smb.conf . Auxiliary parameters can be used to override the default SMB server configuration, but such changes may adversely affect SMB server stability or behavior.

Mounting SMB Share on another machine.

Linux

Verify that your Linux distribution has the required CIFS packages installed. Create a mount point: `sudo mkdir /mnt/smb_share`.

Mount the volume. `sudo mount -t cifs //computer_name/share_name /mnt/smb_share`.

If your share requires user credentials, add the switch `-o username=` with your username after `cifs` and before the share address.

Windows

To mount the SMB share to a drive letter on Windows, open the command line and run the following command with the appropriate drive letter, computer name, and share name.

```
net use Z: \\computer_name\share_name /PERSISTENT:YES
```

Apple

Open **Finder > Go > Connect To Server** Enter the SMB address: `smb://192.168.1.111`.

Input the username and password for the user assigned to that pool or guest if the share has guest access.

FreeBSD

Create a mount point: `sudo mkdir /mnt/smb_share`.

Mount the volume. `sudo mount_smbfs -I computer_name\share_name /mnt/smb_share`.

SMB Home Shares

TrueNAS offers the **Use as Home Share** option for organizations or SMEs that want to use a single SMB share to provide a personal directory to every user account.

The **Use as Home Share** feature is available for a single TrueNAS SMB share. You can create additional SMB shares without the **Use as Home Share** option enabled.

Create a Pool and Join Active Directory

First, go to **Storage** and [create a pool](#).

Next, [set up the Active Directory](#) that you want to share resources with over your network.

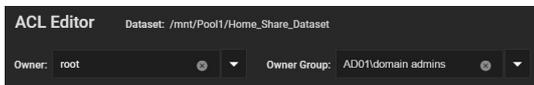
Prepare a Dataset

Go to **Storage** and open the **Storage** page next to the root dataset in the pool you just created, then click **Add Dataset**.

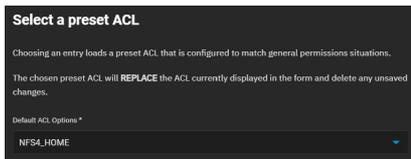
Name the dataset and set the **Share Type** to **SMB**.

After creating the dataset, go to **Storage** and open the **Storage** page next to the new dataset. Select **View Permissions**, then click

Click the **Group** drop-down list and change the owning group to your Active Directory's domain admins.



Click **Use an ACL Preset** and choose **NFS4_HOME**. Then, click **Continue**.



Create the Share

Go to **Shares > Windows (SMB) Shares** and click **Add**.

Set the **Path** to the prepared dataset.

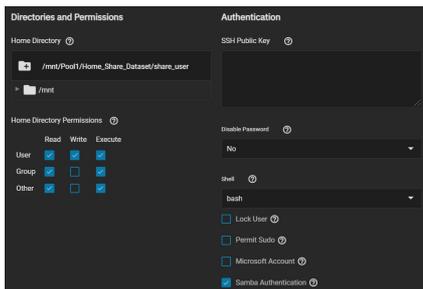
The **Name** automatically becomes identical to the dataset. Leave this at the default.

Set the **Purpose** to **No presets**, then click **Advanced Options** and set **Use as Home Share**. Click **Save**.

Enable the **SMB** service in **System Settings > Services** to make the share is available on your network.

Add Users

Go to **Credentials > Local Users** and click **Add**. Create a new user name and password. By default, the user **Home Directory** is titled from the user account name and added as a new subdirectory of **Home_Share_Dataset**.



If existing users require access to the home share, go to **Credentials > Local Users** and edit an existing account.

Adjust the user's home directory to the appropriate dataset and give it a name to create their own directory.

After adding the user accounts and configuring permissions, users can log in to the share and see a folder matching their username.

SMB Shadow Copies

[Shadow Copies](#), also known as the Volume Shadow Copy Service (VSS) or Previous Versions, is a Microsoft service for creating volume snapshots. You can use shadow copies to restore previous versions of files from within Windows Explorer.

By default, all ZFS snapshots for a dataset underlying an SMB share path are presented to SMB clients through the volume shadow copy service or are accessible directly with SMB when the hidden ZFS snapshot directory is within the SMB share's path.

There are a few caveats about shadow copies to be aware of before activating the feature in TrueNAS:

- Shadow copies might not work if the Windows system is not patched to the latest service pack. If no previous versions of files to restore are visible, use Windows Update to ensure the system is fully up-to-date.
- Shadow copy support only works for ZFS pools or datasets.
- Permissions on the pool or dataset SMB shares must be configured appropriately.
- Shadow copy cannot be deleted by users with an SMB client. Instead, the administrator uses the TrueNAS web interface to remove snapshots. Shadow copies can be disabled for an SMB share by unsetting **Enable shadow copies** for the SMB share. This does not prevent access to the hidden .zfs/snapshot directory for a ZFS dataset when the directory is located within the path for an SMB share.

To enable shadow copies, go to **Shares > Windows (SMB) Shares** and **Edit** an existing share. Open the **Advanced Options** and set **Enable Shadow Copies**.

Windows 10 v2004 Issue

Some users have experienced issues in the Windows 10 v2004 release where network shares can't be accessed. The problem appears to come from a bug in gpedit.msc, the Local Group Policy Editor. Unfortunately, setting the **Allow insecure guest logon** flag value to **Enabled** in **Computer Configuration > Administrative Templates > Network > Lanman Workstation** appears to have no effect on the configuration.

To work around this issue, edit the Windows registry. Use **Regedit** and go to **HKLM\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters**. The **DWORD AllowInsecureGuestAuth** is an incorrect value: **0x00000000**. Change this value to **0x00000001** (Hexadecimal 1) to allow adjusting the settings in gpedit.msc. You can use a Group Policy Update to apply this to a fleet of Windows machines.

UNIX (NFS) Shares

Creating a Network File System (NFS) share on TrueNAS makes a lot of data available for anyone with share access. Depending on the share's configuration, it can restrict users to read or write privileges.

To create a new share, make sure a dataset is available with all the data for sharing.

Creating an NFS Share



Video Player is loading.
 Video URL: <https://www.udenas.com/docs/files/scaleangelfishnfsshare.mp4>

Play Video

Play

Mute

Go to **Shares > Unix (NFS) Shares** and click **Add**.

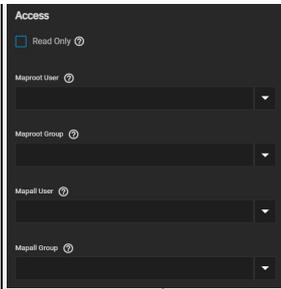
- captions off, selected

Use the file browser to select the dataset to be shared. You can enter an optional text to help identify the share in **Description**. Clicking **Save** creates the share. At the time of creation, you can select **Enable Service** for the service to start and to automatically start after any reboots. If you wish to create the share but not immediately enable it, select **Cancel**.

- Unknown, selected

NFS Share Settings

Setting	Description
Path	Type or browse to the full path to the pool or dataset to share. Click Add to configure multiple paths.
Description	Enter a brief description of the share. Click Save to create the share and cancel and close the window.
Enabled	Enable this NFS share. Unset to disable this NFS share without deleting the configuration.
Add networks	Enter an allowed network in CIDR notation. Click Add to define another authorized network. Defining an authorized network restricts access to all other networks. Leave empty to allow all networks.
Add hosts	Enter a hostname or IP address to allow that system access to the NFS share. Click Add to define another allowed system. Defining authorized systems restricts access to all other systems. Leave the field empty to allow all systems access to the share.
Advanced Options	Window
	Color: Black Transparency: Transparent
	Opening the Advanced Options allows tuning the share access permissions and defining authorized networks.
	Font Size
	100%



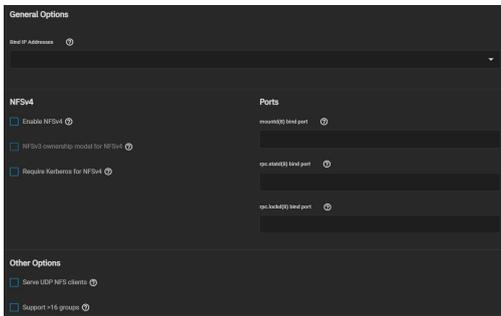
Setting	Value	Description
Read Only	checkbox	Prohibits writing to the share when set.
Maproot User	string or drop-down	Select a user to apply that user's permissions to the <i>root</i> user.
Maproot Group	string or drop-down	Select a group to apply that group's permissions to the <i>root</i> user.
Mapall User	string or drop-down	Permissions for the chosen user applied to all clients.
Mapall Group	string or drop-down	Permissions for the chosen group are applied to all clients.

To edit an existing NFS share, go to **Shares > Unix Shares (NFS)** and click the share you want to edit. The options available are identical to the share creation options.

Configure the NFS Service

To begin sharing, go to **System Settings > Services** and click the **NFS** toggle to running. Select **Start Automatically** if you want NFS to activate when TrueNAS boots.

NFS service settings can be configured by clicking



Setting	Description
Bind IP Addresses	Select IP addresses to listen to for NFS requests. Leave empty for NFS to listen to all available addresses.
Enable NFSv4	Set to switch from NFSv3 to NFSv4.
NFSv3 ownership model for NFSv4	Set when NFSv4 ACL support is needed without requiring the client and the server to sync users and groups.
Require Kerberos for NFSv4	Set to force NFS shares to fail if the Kerberos ticket is unavailable.
Serve UDP NFS clients	Set if NFS clients need to use the User Datagram Protocol (UDP).
Support >16 groups	Set when a user is a member of more than 16 groups. This assumes group membership is configured correctly on the NFS server.
mountd(8) bind port	Enter a number to bind mountd only to that port.
rpc.statd(8) bind port	Enter a number to bind rpc.statd only to that port.
rpc.lockd(8) bind port	Enter a number to bind rpc.lockd only to that port.

Unless you need a specific setting, we recommend using the default NFS settings.

When TrueNAS is already connected to [Active Directory](#), setting **NFSv4** and **Require Kerberos for NFSv4** also requires a [Kerberos Keytab](#).

Connecting to the NFS Share

Although you can connect to an NFS share with various operating systems, it is recommended to use a Linux/Unix operating system. First, download the `nfs-common` kernel module. This can be done using the installed distribution's package manager. For example, on Ubuntu/Debian, enter `sudo apt-get install nfs-common` in the terminal.

After installing the module, connect to an NFS share by entering `sudo mount -t nfs {IPaddressOfTrueNASsystem}:{path/to/nfsShare} {localMountPoint}`. In the above example, `{IPaddressOfTrueNASsystem}` is the remote TrueNAS system's IP address that contains the NFS share, `{path/to/nfsShare}` is the path to the NFS share on the TrueNAS system, and `{localMountPoint}` is a local directory on the host system configured for the mounted NFS share. For example, `sudo mount -t nfs 10.239.15.110:/mnt/Pool1/NFS_Share /mnt` mounts the NFS share **NFS_Share** to the local directory `/mnt`.

By default, anyone that connects to the NFS share only has read permission. To change the default permissions, edit the share, open the **Advanced Options**, and change the **Access** settings.

ESXi 6.7 or later is required for read/write functionality with NFSv4 shares.

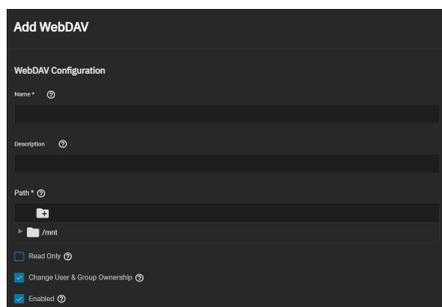
WebDAV

A Web-based Distributed Authoring and Versioning (WebDAV) share makes it easy to share a TrueNAS dataset and its contents over the web.

To create a new share, make sure a dataset is available with all the data for sharing.

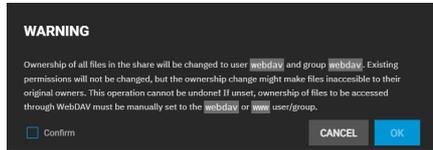
Share Configuration

Go to **Shares > WebDAV Shares** and click **Add**.



Enter a share **Name** and use the file browser to select the dataset to be shared. An optional **Description** helps to identify the share. To prevent user accounts from modifying the shared data, set **Read Only**.

By default, **Change User & Group Ownership** is set. This changes existing ownership of *all* files in the share to the **webdav** user and group accounts. The default simplifies WebDAV share permission, but is unexpected, so the web interface shows a warning:

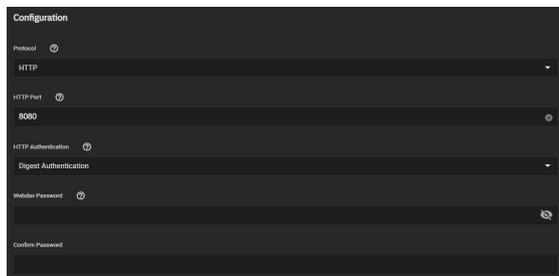


This warning does not show when the **Change User & Group Ownership** checkbox is cleared. In that situation, you must manually set shared file ownership to the **webdav** or **www** user and group accounts.

By default, the new WebDAV share is immediately active. To create the share but not immediately activate it, unset **Enable**. Click **Submit** to create the share.

Service Activation

Creating a share allows users to activate the WebDAV service. To enable or disable the WebDAV service, go to **System Settings > Services** and toggle **WebDAV**. To automatically start the service when TrueNAS boots, set **Start Automatically**. Click **Save** to change the service settings.



For better data security, set the **Protocol** to **HTTPS**. If you require it, you must choose an SSL certificate (*freenas_default* is always available). All **Protocol** options require you to define a number in the **Port** field. Make sure the network is not already using the WebDAV Service port.

To prevent unauthorized access to the shared data, set the **HTTP Authentication** to either **Basic** or **Digest** and create a new **Webdav Password**.

Be sure to click **Save** after making any changes.

Connecting to the WebDAV Share

WebDAV shared data is accessible from a web browser. To see the shared data, open a new browser tab and enter `{PROTOCOL}://{TRUENASIP}:{PORT}/{SHAREPATH}`. Replace the elements in curly brackets `{}` with your chosen WebDAV share and service settings. Example: `https://10.2.1.1:8081/newdataset`

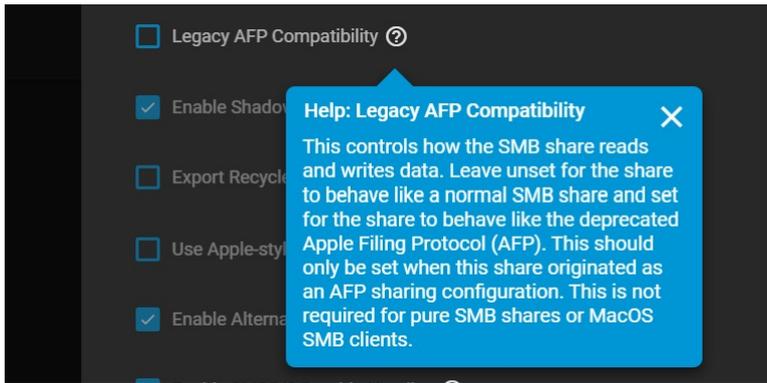
TrueNAS requires a username and password when setting the **Authentication** WebDAV service option to **Basic** or **Digest**. Enter the user name **webdav** and the password defined in the WebDAV service.

5.1 - AFP Migration

Since the Apple Filing Protocol (AFP) for shares has been deprecated and no longer receives updates, it is not included in TrueNAS SCALE.

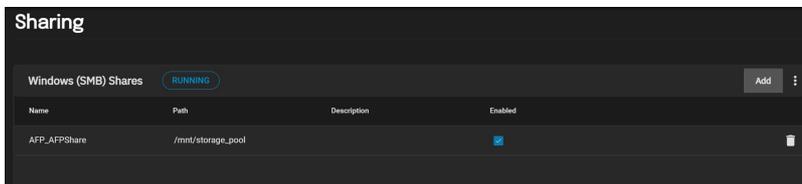
However, since users can sidegrade a TrueNAS CORE configuration into SCALE, TrueNAS SCALE will migrate any previously saved AFP configurations into SMB configurations.

The Windows (SMB) Shares advanced options section has a checkbox that enables compatibility for shares that were previously AFP. The check box must be set when the share is AFP-related or the data can become corrupted. Do **not** set the checkbox if you want a share to be pure SMB (no AFP relation).



As of SCALE version 21.06, the Netatalk service has been removed. AFP shares will be automatically migrated to SMB shares with the Legacy AFP Compatibility box checked. Do not uncheck Legacy AFP Compatibility as it will impact how data is written to and read from shares. Any other shares created to access these paths after the migration *must* also have the Legacy AFP Compatibility box checked.

Once you have [sidegraded from CORE to SCALE](#), you can find your migrated AFP configuration in **Shares > Windows Shares (SMB)** with the prefix AFP_. To make the migrated AFP share accessible, start the SMB service.



6 - Data Protection

The Data Protection section allows users to set up multiple redundant tasks that will protect and/or backup data in case of drive failure.

Scrub Tasks and S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) Tests can provide early disk failure alerts by identifying data integrity problems and detecting various indicators of drive reliability.

Cloud Sync, Periodic Snapshot, Rsync, and Replication Tasks, provide backup storage for data and allow users to revert the system to a previous configuration or point in time.

Ready to get started? Choose a topic or article from the left-side Navigation pane. Click the < symbol to expand the menu to show the topics under this section.

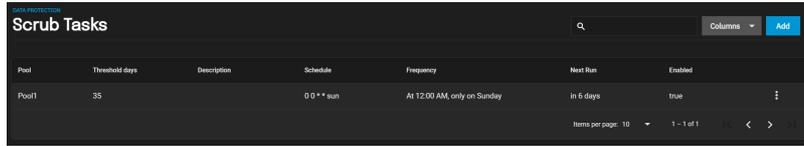
6.1 - Scrub Tasks

- [Default Scrub Tasks](#)
- [Creating New Scrub Tasks](#)
- [Editing Scrub Tasks](#)

When TrueNAS performs a scrub, ZFS scans the data on a pool. Scrubs identify data integrity problems, detect silent data corruptions caused by transient hardware issues, and provide early disk failure alerts.

Default Scrub Tasks

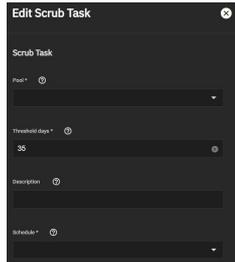
TrueNAS generates a default scrub task when you create a new pool and sets it to run every Sunday at 12:00 AM.



Creating New Scrub Tasks

TrueNAS needs at least one data [pool](#) to create scrub task.

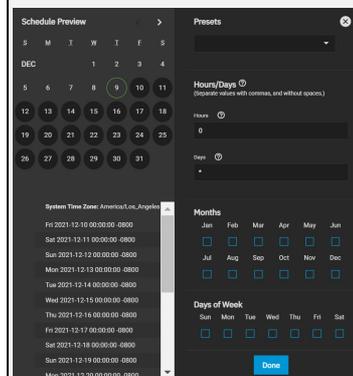
To create a scrub task for a pool, go to **Data Protection** and click **ADD** in the **Scrub Tasks** window.



Scrub Task

Name	Description
Pool	Choose a pool to scrub.
Threshold days	Controls the task schedule by setting how many days must pass before a completed scrub can run again. If you schedule a scrub to run daily and set <i>Threshold days</i> to 7, the scrub attempts to run daily. If the scrub succeeds, it will check but won't run again until seven days pass. Using a multiple of seven ensures the scrub runs on the same weekday.
Description	Describe the scrub task.
Schedule	How often to run the scrub task. Choose one of the presets or <i>Custom</i> to use the Advanced Scheduler .
Enabled	Unset to disable the scheduled scrub without deleting it.

Advanced Scheduler



Choosing a **Presets** option populates the rest of the fields. To customize a schedule, enter [cron](#) values for the Minutes/Hours/Days.

These fields accept standard [cron](#) values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering *10* means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means match all values.

You can set specific time ranges by entering hyphenated number values. For example, entering *30-35* in the **Minutes** field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

You can also enter lists of values. Enter individual values separated by a comma (,). For example, entering *1,14* in the **Hours** field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, entering * in **Days** runs the task every day of the month. Entering */2 runs it every other day.

Combining the above examples creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

TrueNAS has an option to select which **Months** the task runs. Leaving each month unset is the same as selecting every month.

The **Days of Week** schedules the task to run on specific days in addition to any listed days. For example, entering *1* in **Days** and setting *Wed* for **Days of Week** creates a schedule that starts a task on the first day of the month and every Wednesday of the month.

The **Schedule Preview** displays when the current settings mean the task runs.

Examples of CRON syntax

Syntax	Meaning	Examples
*	Every item.	* (minutes) = every minute of the hour. * (days) = every day.
*/N	Every N th item.	*/15 (minutes) = every 15th minute of the hour. */3 (days) = every 3rd day. */3 (months) = every 3rd month.

Comma and hyphen/dash	Each stated item (comma) Each item in a range (hyphen/dash).	1,31 (minutes) = on the 1st and 31st minute of the hour. 1-3,31 (minutes) = on the 1st to 3rd minutes inclusive, and the 31st minute, of the hour. mon-fri (days) = every Monday to Friday inclusive (every weekday). mar,jun,sep,dec (months) = every March, June, September, December.
-----------------------	---	---

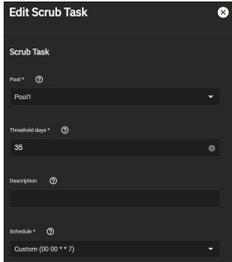
You can specify days of the month or days of the week.

TrueNAS lets users create flexible schedules using the available options. The table below has some examples:

Desired schedule	Values to enter
3 times a day (at midnight, 08:00 and 16:00)	months=*; days=*; hours=0/8 or 0,8,16; minutes=0 (Meaning: every day of every month, when hours=0/8/16 and minutes=0)
Every Monday/Wednesday/Friday, at 8.30 pm	months=*; days=mon,wed,fri; hours=20; minutes=30
1st and 15th day of the month, during October to June, at 00:01 am	months=oct-dec,jan-jun; days=1,15; hours=0; minutes=1
Every 15 minutes during the working week, which is 8am - 7pm (08:00 - 19:00) Monday to Friday	Note that this requires two tasks to achieve: (1) months=*; days=mon-fri; hours=8-18; minutes=*/15 (2) months=*; days=mon-fri; hours=19; minutes=0 We need the second scheduled item, to execute at 19:00, otherwise we would stop at 18:45. Another workaround would be to stop at 18:45 or 19:45 rather than 19:00.

Editing Scrub Tasks

To edit a scrub, go to **Data Protection** and click the scrub task you want to edit.



6.2 - Cloud Sync Tasks

- [Requirements](#)
- [Creating a Cloud Sync Task](#)
 - [Testing Settings](#)
- [Cloud Sync Behavior](#)
- [Cloud Sync Restore](#)

TrueNAS can send, receive, or synchronize data with a cloud storage provider. Cloud sync tasks allow for single-time transfers or recurring transfers on a schedule. They are an effective method to back up data to a remote location.

Using the cloud means data can go to a third-party commercial vendor not directly affiliated with iXsystems. You should fully understand vendor pricing policies and services before using them for cloud sync tasks.

iXsystems is not responsible for any charges incurred from using third-party vendors with the cloud sync feature.

TrueNAS supports major providers like Amazon S3, Google Cloud, and Microsoft Azure. It also supports many other vendors. To see the full list of supported vendors, go to **Credentials > Backup Credentials > Cloud Credentials** click **Add** and open the **Provider** drop-down list.

Requirements

- You must have all system [storage](#) configured and ready to receive or send data.
- You must have a cloud storage provider account and location (like an Amazon S3 bucket).
- You must have cloud storage account credentials saved to **System > Cloud Credentials** before creating the sync task. See [Cloud Credentials](#) for specific instructions.

Creating a Cloud Sync Task

Cloud Synch Tutorial Video



Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaleangelfishcloudsync.mp4>

Play Video

Go to **Data Protection > Cloud Sync Tasks** and click **Add**.

Play

Mute

Current Time 0:00

/

Duration 1:13

Loaded: 100.00%

Stream Type LIVE

Seek to live, currently behind liveLIVE

Remaining Time -1:13

1x

Playback Rate

Chapters

the window.

Type a memorable task description in the **Description** field. Use the **Credential** dropdown list to select an existing cloud or create a new one with the **+ Add a backup credential** option. TrueNAS connects to the chosen cloud storage provider and shows the available storage locations. Select the option if data is transferring to (**PUSH**) or from (**PULL**) the cloud storage location (**Remote**). Select a **Transfer Mode**:

Sync

Sync keeps all the files identical between the two storage locations. If the sync encounters an error, it does not delete files in the destination. One common error occurs when the [Dropbox copyright detector](#) flags a file as copyrighted.

Note that syncing to a Backblaze B2 bucket does not delete files from the bucket, even when you deleted those files locally. Instead, files are tagged with a version number or moved to a hidden state. To automatically delete old or unwanted files from the bucket, adjust the [Backblaze B2 Lifecycle Rules](#).

Sync cannot delete files stored in Amazon S3 Glacier or S3 Glacier Deep Archive. These files must first be restored by another means, like the [Amazon S3 console](#).

Copy

Copy duplicates each source file into the destination, overwriting any destination files using the same name as the source. Copying is the least potentially destructive option.

Move

Move transfers the files from the source to the destination and deletes the source files. **Move** also overwrites files with the same names on the destination.

Next, use the **Control** options. Define when the task runs using **Schedule**. If you need a specific schedule, select **Custom** and use the **Advanced Scheduler**.

Choosing a **Presets** option populates the rest of the fields. To customize a schedule, enter [cron](#) values for the Minutes/Hours/Days.

These fields accept standard [cron](#) values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering *10* means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means match all values.

You can set specific time ranges by entering hyphenated number values. For example, entering *30-35* in the **Minutes** field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

You can also enter lists of values. Enter individual values separated by a comma (,). For example, entering *1,14* in the **Hours** field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, entering * in **Days** runs the task every day of the month. Entering */2 runs it every other day.

Combining the above examples creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

TrueNAS has an option to select which **Months** the task runs. Leaving each month unset is the same as selecting every month.

The **Days of Week** schedules the task to run on specific days in addition to any listed days. For example, entering *1* in **Days** and setting *Wed* for **Days of Week** creates a schedule that starts a task on the first day of the month *and* every Wednesday of the month.

The **Schedule Preview** displays when the current settings mean the task runs.

Examples of CRON syntax

Syntax	Meaning	Examples
*	Every item.	* (minutes) = every minute of the hour. * (days) = every day.
*/N	Every N th item.	*/15 (minutes) = every 15th minute of the hour. */3 (days) = every 3rd day. */3 (months) = every 3rd month.
Comma and hyphen/dash	Each stated item (comma) Each item in a range (hyphen/dash).	1,31 (minutes) = on the 1st and 31st minute of the hour. 1-3,31 (minutes) = on the 1st to 3rd minutes inclusive, and the 31st minute, of the hour. mon-fri (days) = every Monday to Friday inclusive (every weekday). mar,jun,sep,dec (months) = every March, June, September, December.

You can specify days of the month or days of the week.

TrueNAS lets users create flexible schedules using the available options. The table below has some examples:

Desired schedule	Values to enter
3 times a day (at midnight, 08:00 and 16:00)	months=*; days=*; hours=0/8 or 0,8,16; minutes=0 (Meaning: every day of every month, when hours=0/8/16 and minutes=0)
Every Monday/Wednesday/Friday, at 8.30 pm	months=*; days=mon,wed,fri; hours=20; minutes=30
1st and 15th day of the month, during October to June, at 00:01 am	months=oct-dec,jan-jun; days=1,15; hours=0; minutes=1
Every 15 minutes during the working week, which is 8am - 7pm (08:00 - 19:00) Monday to Friday	Note that this requires two tasks to achieve: (1) months=*; days=mon-fri; hours=8-18; minutes=*/15 (2) months=*; days=mon-fri; hours=19; minutes=0 We need the second scheduled item, to execute at 19:00, otherwise we would stop at 18:45. Another workaround would be to stop at 18:45 or 19:45 rather than 19:00.

Clear the **Enable** checkbox to make the configuration available without allowing the specified schedule to run the task. To manually activate a saved task, go to **Data Protection > Cloud Sync Tasks**, click ► for the cloud sync task you want to run. You are prompted to **CONTINUE** or **CANCEL** the **Run Now** operation.

The remaining options allow tuning the task to your specific requirements.

Specific Options

Transfer

Name	Description
Description	Enter a description of the Cloud Sync Task.
Direction	PUSH sends data to cloud storage. PULL receives data from cloud storage. Changing the direction resets the Transfer Mode to COPY.
Transfer Mode	SYNC: Files on the destination are changed to match those on the source. If a file does not exist on the source, it is also deleted from the destination. COPY: Files from the source are copied to the destination. If files with the same names are present on the destination, they are overwritten. MOVE: After files are copied from the source to the destination, they are deleted from the source. Files with the same names on the destination are overwritten.
Directory/Files	Select the directories or files to be sent to the cloud for Push syncs, or the destination to be written for Pull syncs. Be cautious about the destination of Pull jobs to avoid overwriting existing files.

Remote

Name	Description
Credential	Select the cloud storage provider credentials from the list of available Cloud Credentials.

Control

Name	Description
Schedule	Select a schedule preset or choose Custom to open the advanced scheduler.
Enabled	Enable this Cloud Sync Task. Unset to disable this Cloud Sync Task without deleting it.

Advanced Options

Name	Description
Follow Symlinks	Follow symlinks and copy the items to which they link.
Pre-Script	Script to execute before running sync.
Post-Script	Script to execute after running sync.
Exclude	List of files and directories to exclude from sync. Separate entries by pressing Enter. Examples of proper syntax used to exclude files/directories are: <ul style="list-style-type: none"> • photos will exclude a file named "photos" • /photos will exclude a file named "photos" from root directory (but not subdirectories) • photos/ will exclude a directory named "photos" • /photos/ will exclude a directory named "photos" from root directory (but not subdirectories). See rclone filtering for more details about the --exclude option.

Advanced Remote Options

Name	Description
Remote	PUSH: Encrypt files before transfer and store the encrypted files on the remote system. Files are encrypted using the Encryption Password and Encryption Salt values. PULL: Decrypt files that are being stored on the remote system before the transfer. Transferring the encrypted files

Encryption	requires entering the same Encryption Password and Encryption Salt that was used to encrypt the files. Additional details about the encryption algorithm and key derivation are available in the rclone crypt File formats documentation .
Transfers	Number of simultaneous file transfers. Enter a number based on the available bandwidth and destination system performance. See rclone -transfers .
Bandwidth limit	A single bandwidth limit or bandwidth limit schedule in rclone format. Separate entries by pressing Enter. Example: 08:00,512 12:00,10MB 13:00,512 18:00,30MB 23:00,off. Units can be specified with the beginning letter: b, k (default), M, or G. See rclone -bwlimit .

Scripting and Environment Variables

Advanced users can write scripts that run immediately before or after the cloud task. The **Post-script** field only runs when the cloud sync task succeeds. You can pass a variety of task environment variables into the **Pre-** and **Post-** script fields:

- CLOUD_SYNC_ID
- CLOUD_SYNC_DESCRIPTION
- CLOUD_SYNC_DIRECTION
- CLOUD_SYNC_TRANSFER_MODE
- CLOUD_SYNC_ENCRYPTION
- CLOUD_SYNC_FILENAME_ENCRYPTION
- CLOUD_SYNC_ENCRYPTION_PASSWORD
- CLOUD_SYNC_ENCRYPTION_SALT
- CLOUD_SYNC_SNAPSHOT

There also are provider-specific variables like CLOUD_SYNC_CLIENT_ID or CLOUD_SYNC_TOKEN or CLOUD_SYNC_CHUNK_SIZE.

Remote storage settings:

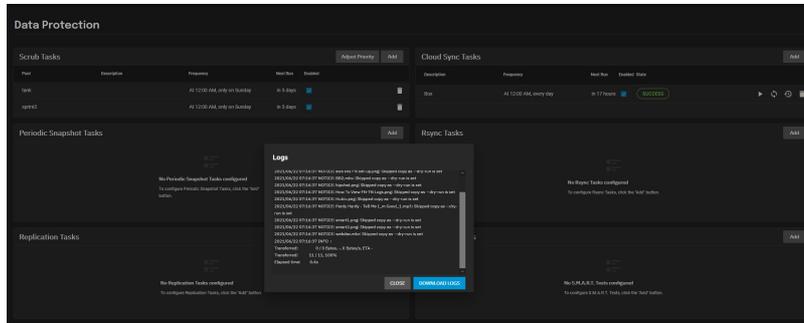
- CLOUD_SYNC_BUCKET
- CLOUD_SYNC_FOLDER

Local storage settings:

- CLOUD_SYNC_PATH

Testing Settings

To test the settings before saving, click **Dry Run**. TrueNAS connects to the cloud storage provider and simulates a file transfer but does not send or receive data. A dialog box displays with the test status and allows you to download the task logs. You can also run this by clicking **Run** after creating the task.



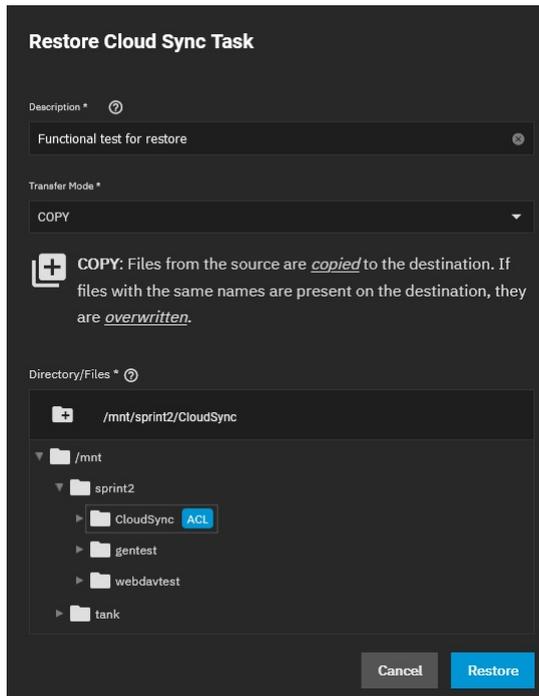
Cloud Sync Behavior

Saved tasks activate according to their schedule or by clicking **Run**. An in-progress cloud sync must finish before another can begin. Stopping an in-progress task cancels the file transfer and requires starting the file transfer over.

To view logs about a running task, or its most recent run, click **State**.

Cloud Sync Restore

To quickly create a new cloud sync that uses the same options but reverses the data transfer, select **Restore** for an existing cloud sync on the **Data Protection** page.



Type a new description for this reversed task, then define the path to a storage location for the transferred data and click **Restore**.

TrueNAS saves the restored cloud sync as another entry in **Data protection > Cloud Sync Tasks**.

If you set the restore destination to the source dataset, TrueNAS may alter ownership of the restored files to **root**. If **root** did not create the original files and

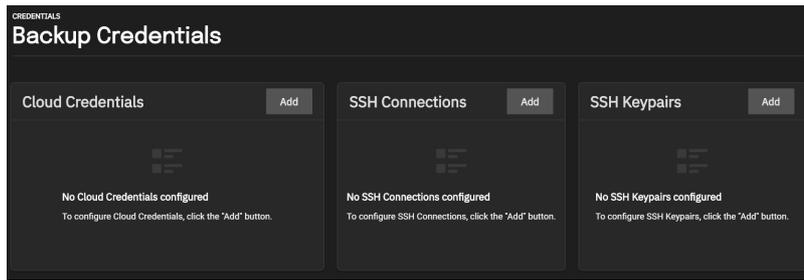
you need them to have a different owner, you can recursively reset their ACL permissions through the GUI or run `chown` from the CLI.

6.2.1 - How To Back Up Google Drive to TrueNAS SCALE

Google Drive and G Suite are widely used tools for creating and sharing documents, spreadsheets, and presentations with team members. While cloud-based tools have inherent backups and replications included by the cloud provider, certain users may require additional backup or archive capabilities. For example, companies using G Suite for important work may be required to keep records for years, potentially beyond the scope of the G Suite subscription. TrueNAS offers the ability to easily back up Google Drive by using the built-in cloud sync.

Setting up Google Drive credentials

Set up the credentials under **Credentials > Backup Credentials**.



Next click **ADD** for *Cloud Credentials*.

Name the Credential and select *Google Drive* for the Provider. Click **LOGIN TO PROVIDER** and login with the appropriate Google user account.

Google will request to allow access to all the Google Drive files for the FreeNAS device.

Authorization

Only proceed if you are setting up cloud sync on your TrueNAS system at <http://192.168.1.123/ui/system/cloudcredentials/add>

[Proceed](#)

Sign in with Google

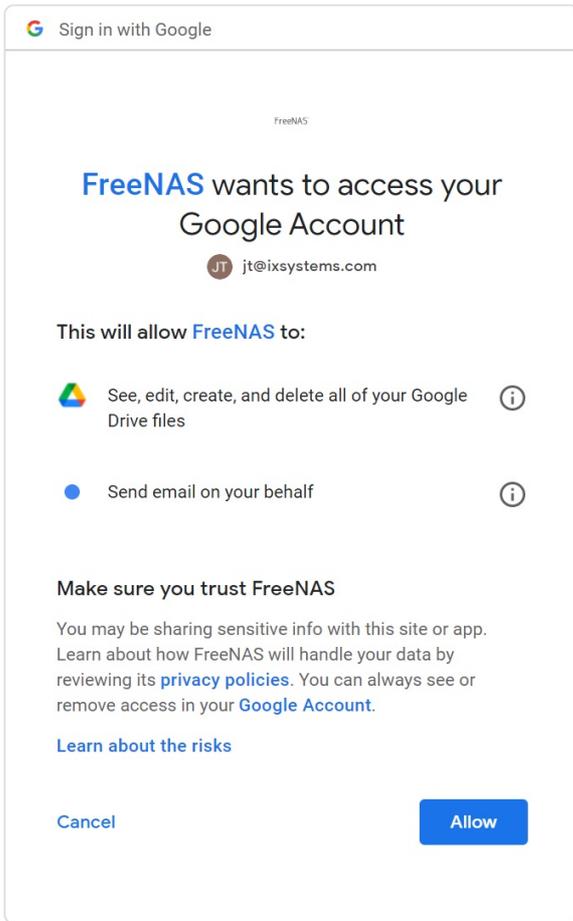
FreeNAS

Choose an account

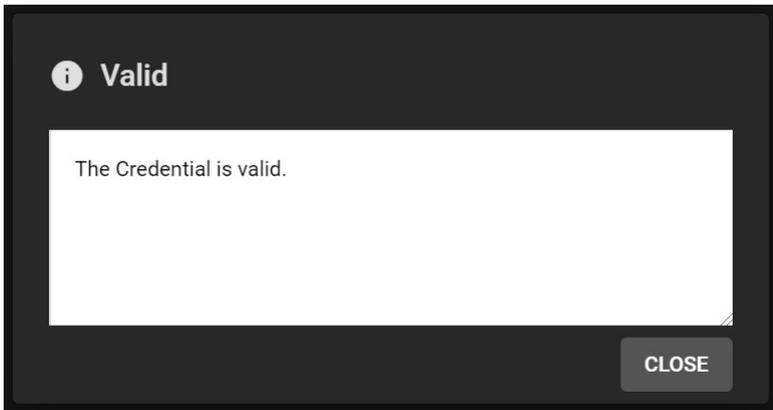
to continue to **FreeNAS**

-  JT Pennington
jt@ixsystems.com
-  Jt Pennington
jt@obs-sec.com
-  Use another account

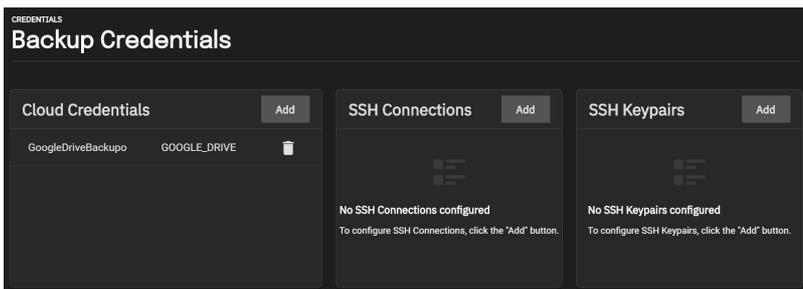
Before using this app, you can review FreeNAS's [privacy policy](#) and terms of service.



Allow access and the appropriate access key will be inserted to the FreeNAS access token. Assign a Team ID if required, but it is not necessary in all cases. Click **VERIFY CREDENTIAL** and wait for the credential to verify.



Once successful, click **SUBMIT**. The new cloud credentials will be visible in the web interface.



Set the cloud sync task

Click on **Data Protection** to open the Tasks page. Click on + icon for *Cloud Sync Tasks* to create a new Cloud Sync Task.

Set the backup time frame, frequency, and folders - both the cloud-based folder and TrueNAS dataset. Set whether the synchronization should sync all changes, just copy new files, or move files. Files are removed from the cloud source or TrueNAS source depending on push or pull. Add a description for the task and select the cloud credentials. Choose the appropriate cloud folder target and TrueNAS storage location.

Select the file transfer mode:

- **Sync:** Keep files newly created or deleted the same.
- **Copy:** Copy new files to the appropriate target (i.e., TrueNAS pulls files from Google Drive or pushes files to Google Drive).
- **Move:** Copy files to the target and then delete files from the source. Using Move, users can set a folder in Google Drive for archival, and move older documents to that folder from their Drive account. Those files would then automatically get backed up to their TrueNAS storage.

Transfer

Description * ?

Direction * ?

Transfer Mode * ?

+ COPY: Files from the source are copied to the destination. If files with the same names are present on the destination, they are overwritten.

Directory/Files * ?

 /mnt
 temp
 Temp **ACL**

Remote

Credential * ?

Folder ?

 /

Control

Schedule ?

Enabled ?

Advanced Options

Follow Symlinks ?

Pre-script ?

Post-script ?

Exclude ?

Advanced Remote Options

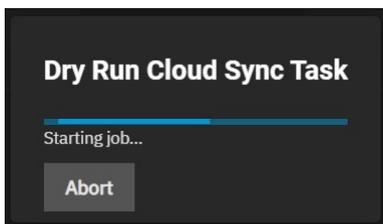
Use --fast-list ?

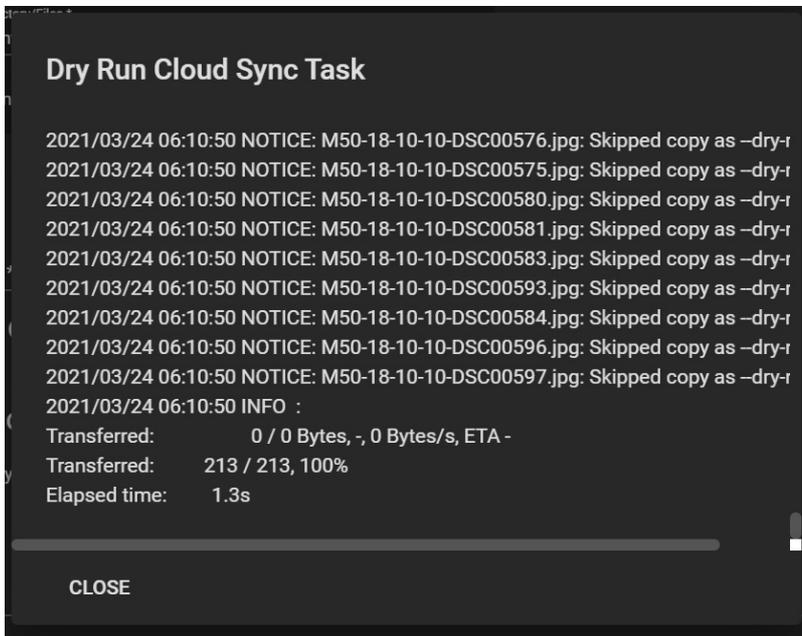
Remote Encryption ?

Transfers ?
 Bandwidth Limit ?

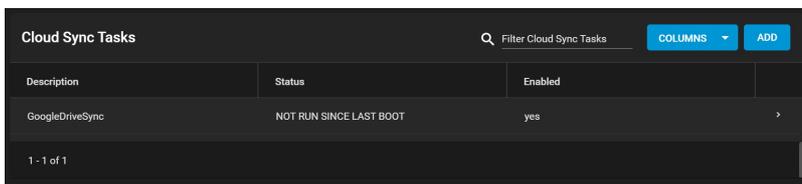
Save Cancel Dry Run

Once created, attempt a Dry Run of the task.

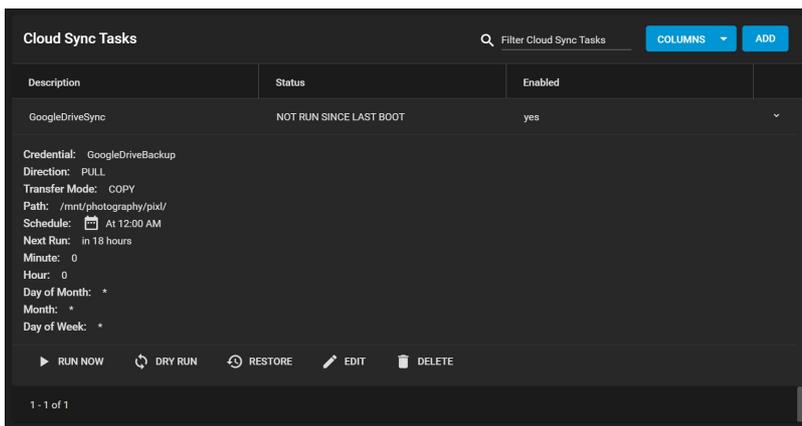




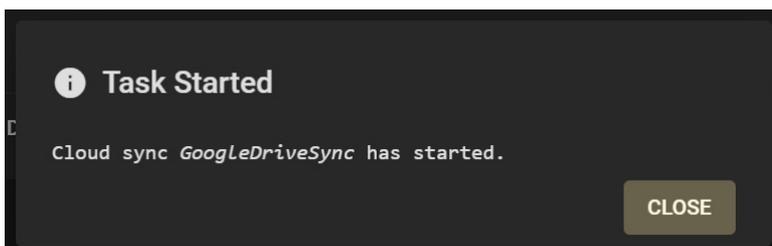
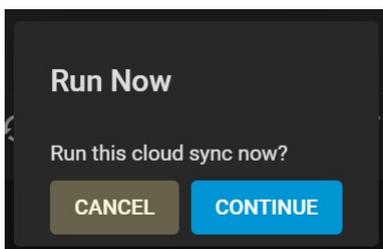
If the Dry Run succeeds, click **SUBMIT** to save the task.

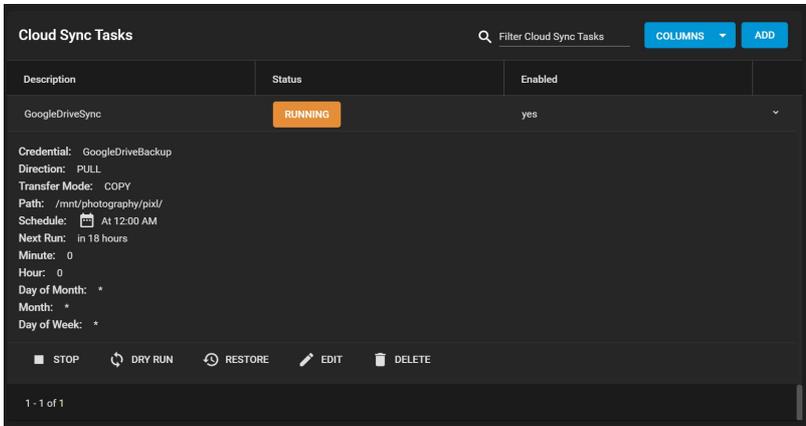


Expand the section down to see the options for the task.



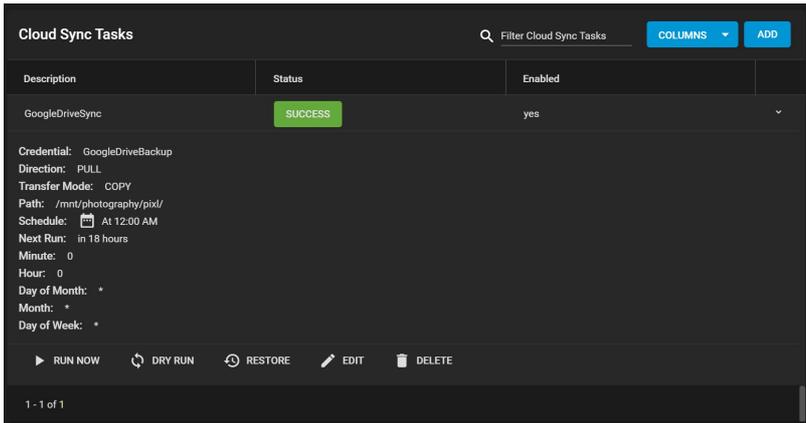
Clicking **RUN NOW** will prompt the task to start immediately.





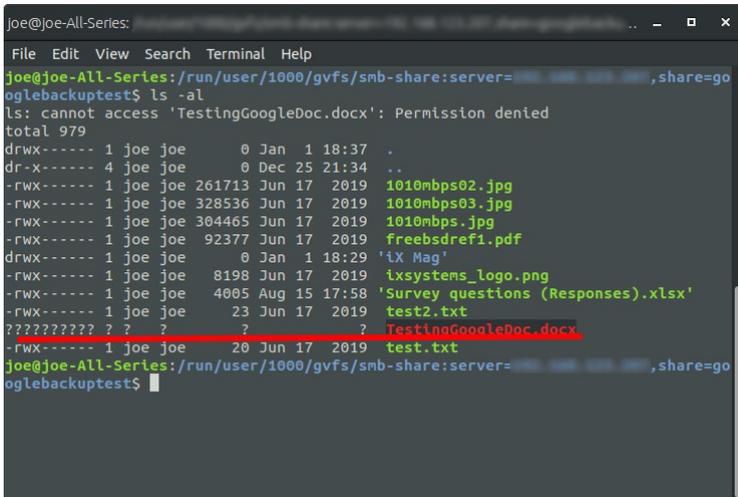
The web interface will show the status as **RUNNING** and **SUCCESS** upon completion. Details can be accessed via the **Task Manager** icon in the upper right-hand corner.

Once the sync reports a status of **SUCCESS** you can verify this by opening the folder on another computer if it is a share, through SSH access, or by checking the destination directory through the TrueNAS CLI.

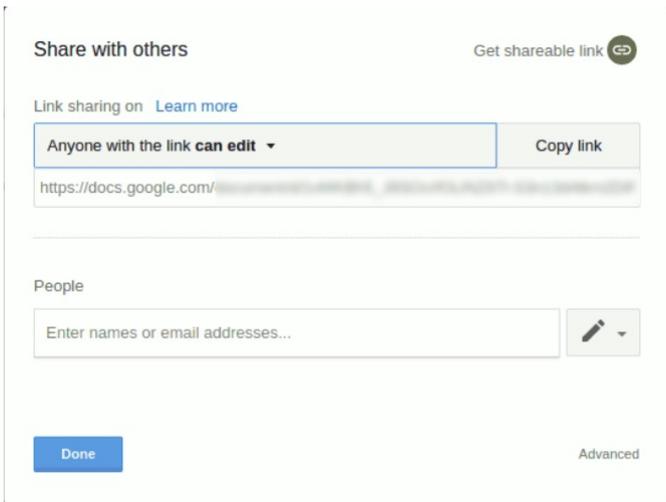


Working with Google created content

One caveat is that Google Docs and other files created with Google tools have their own proprietary set of permissions and their read/write characteristics unknown to the system over a standard file share. Files are unreadable as a result.



To allow Google created files to become readable, allow link sharing to access the files before the backup. Doing so ensures that other users can open the files with read access, make changes, and then save as another file if further edits are needed. Note that this is only necessary if the file was created using Google Docs, Google Sheets, or Google Slides; other files should not require modification of their share settings.



TrueNAS is perfect for storing content, including cloud-based content, for the long-term. Not only is it simple to sync and backup from the cloud, but users can rest assured that their data is safe, with snapshots, copy-on-write, and built-in replication functionality.

6.3 - Rsync Tasks

- [Basic Requirements](#)
- [Creating an Rsync Task](#)
- [Rsync Service and Modules](#)

You often need to copy data to another system for backup or when migrating to a new system. A fast and secure way of doing this is by using [rsync](#). These instructions assume that both sides of the rsync task, host and remote, use a TrueNAS systems.

Basic Requirements

Rsync requires a [dataset](#) with the needed data on either the host or remote system. Rsync provides the ability to either push or pull data. When using the **Rsync Tasks** function to push, data is copied from a host system to a remote system. When using the **Rsync Tasks** function to pull, data is pulled from a remote system and put on the host system.

The remote system must have the rsync service activated. Additional requirements are listed further down for either rsync module and SSH tasks.

Creating an Rsync Task

Tutorial Video



Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaleangelfishrsync.mp4>

Play Video

Play

Mute

Go to **Data Protection** > **Rsync Tasks** and click **Add**. The **Add Rsync Task** configuration screen displays.

Current Time 0:00

/

Duration 0:43

Loaded: 100.00%

Stream Type LIVE

Seek to live, currently behind liveLIVE

Remaining Time -0:43

1x

Playback Rate

Chapters

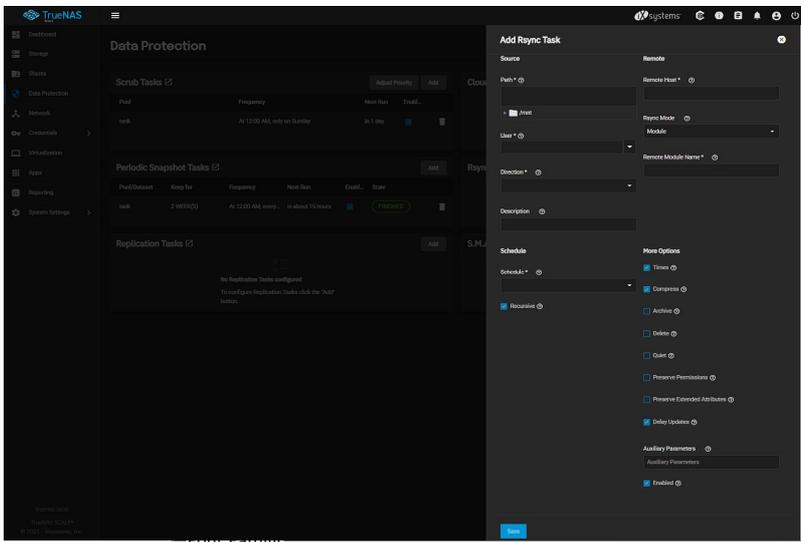
- Chapters

Descriptions

- descriptions off, selected

Captions

- captions off, selected



The **Rsync Mode** field has two primary rsync modes: **Module** and **SSH**. Each mode has different requirements. See the related tabs below for each rsync mode.

Module

Module Requirements

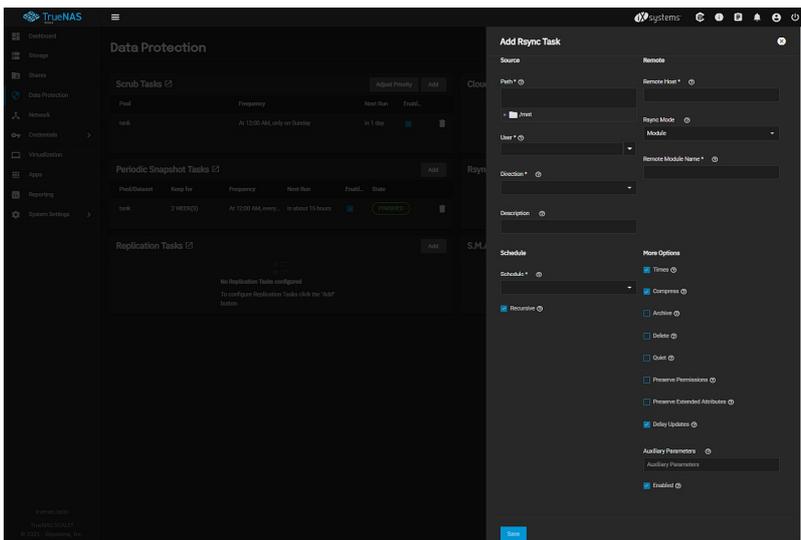
Before you create an rsync task on the host system, you must create a module on the remote system.

When TrueNAS is the remote system, create a module. Go to **System > Services** and click instructions in the [Rsync Service section](#) of this article.

for the **Rsync** service. See the specific configuration

Rsync Module Process

Log in to the host system interface, go to **Data Protection > Rsync Tasks**, and click **Add**. The **Add Rsync Task** configuration panel displays. Use the scroll bar displayed at the right to access all configuration fields.



Select the source dataset to use with the rsync task and a user account to run the rsync task. Choose a direction for the rsync task.

Select a schedule for the rsync task. If you need a custom schedule, select **Custom**.

Advanced Scheduler

Choosing a **Presets** option populates the rest of the fields. To customize a schedule, enter [cron](#) values for the Minutes/Hours/Days.

These fields accept standard [cron](#) values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering **10** means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means match all values.

You can set specific time ranges by entering hyphenated number values. For example, entering **30-35** in the **Minutes** field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

You can also enter lists of values. Enter individual values separated by a comma (,). For example, entering **1,14** in the **Hours** field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, entering * in **Days** runs the task every day of the month. Entering */2 runs it every other day.

Combining the above examples creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

TrueNAS has an option to select which **Months** the task runs. Leaving each month unset is the same as selecting every month.

The **Days of Week** schedules the task to run on specific days in addition to any listed days. For example, entering 1 in **Days** and setting *Wed* for **Days of Week** creates a schedule that starts a task on the first day of the month *and* every Wednesday of the month.

The **Schedule Preview** displays when the current settings mean the task runs.

Examples of CRON syntax

Syntax	Meaning	Examples
*	Every item.	* (minutes) = every minute of the hour. * (days) = every day.
*/N	Every N th item.	*/15 (minutes) = every 15th minute of the hour. */3 (days) = every 3rd day. */3 (months) = every 3rd month.
Comma and hyphen/dash	Each stated item (comma) Each item in a range (hyphen/dash).	1,31 (minutes) = on the 1st and 31st minute of the hour. 1-3,31 (minutes) = on the 1st to 3rd minutes inclusive, and the 31st minute, of the hour. mon-fri (days) = every Monday to Friday inclusive (every weekday). mar,jun,sep,dec (months) = every March, June, September, December.

You can specify days of the month or days of the week.

TrueNAS lets users create flexible schedules using the available options. The table below has some examples:

Desired schedule	Values to enter
3 times a day (at midnight, 08:00 and 16:00)	months=*; days=*; hours=0/8 or 0,8,16; minutes=0 (Meaning: every day of every month, when hours=0/8/16 and minutes=0)
Every Monday/Wednesday/Friday, at 8.30 pm	months=*; days=mon,wed,fri; hours=20; minutes=30
1st and 15th day of the month, during October to June, at 00:01 am	months=oct-dec,jan-jun; days=1,15; hours=0; minutes=1
Every 15 minutes during the working week, which is 8am - 7pm (08:00 - 19:00) Monday to Friday	Note that this requires two tasks to achieve: (1) months=*; days=mon-fri; hours=8-18; minutes=*/15 (2) months=*; days=mon-fri; hours=19; minutes=0 We need the second scheduled item, to execute at 19:00, otherwise we would stop at 18:45. Another workaround would be to stop at 18:45 or 19:45 rather than 19:00.

Next, enter the **Remote Host** IP address or hostname. Use the format *username@remote_host* when the username differs from the host entered into the **Remote Host** field. Select **Module** in the **Rsync Mode** drop-down. Type the **Remote Module Name** exactly as it appears on the remote system.

Configure the remaining options according to your specific needs.

Options

Source

Name	Description
Path	Browse to the path to be copied. FreeBSD file path limits apply. Other operating systems can have different limits which might affect how they can be used as sources or destinations.
User	Select the user to run the rsync task. The user selected must have permissions to write to the specified directory on the remote host.
Direction	Direct the flow of data to the remote host. During a <i>push</i> , the dataset transfers to the remote module. During a <i>pull</i> , the dataset stores files from the <i>remote</i> system.
Description	Enter a description of the rsync task.

Schedule

Name	Description
Schedule	Select a schedule preset or choose Custom to open the advanced scheduler.
Recursive	Set to include all subdirectories of the specified directory. When unset, only the specified directory is included.

Remote

Name	Description
Remote Host	Enter the IP address or hostname of the remote system that will store the copy. Use the format <i>username@remote_host</i> if the username differs on the remote host.
Rsync Mode	Choose to either use a custom-defined remote module of the rsync server or to use an SSH configuration for the rsync task.

More Options

Name	Description
Times	Set to preserve modification times of files.
Compress	Set to reduce the size of data to transmit. Recommended for slow connections.
Archive	When set, rsync is run recursively, preserving symlinks, permissions, modification times, group, and special files. When run as root, owner, device files, and special files are also preserved. Equivalent to passing the flags <code>-rLptgo</code> to rsync.
Delete	Delete files in the destination directory that do not exist in the source directory.
Quiet	Set to suppress informational messages from the remote server.
Preserve Permissions	Set to preserve original file permissions. This is useful when the user is set to root.
Preserve Extended Attributes	Extended attributes are preserved, but must be supported by both systems.
Delay Updates	Set to save the temporary file from each updated file to a holding directory until the end of the transfer when all transferred files are renamed into place.
Auxiliary Parameters	Additional rsync(1) options to include. Separate entries by pressing Enter. Note: The <code>"</code> character must be escaped with a backslash (<code>\"</code>).txt) or used inside single quotes (<code>'*.txt'</code>).
Enabled	Enable this rsync task. Unset to disable this rsync task without deleting it.

The **Module** mode adds the **Remote Module Name** field to the **Remote** section. You must define at least one module in [rsyncd.conf\(5\)](#) of the rsync server or in the rsync modules of another system.

If the **Enable** checkbox is not selected it disables the task schedule, but you can still save the rsync task and run it manually.

SSH

SSH Requirements

The remote system must have SSH enabled. To enable SSH in TrueNAS, go to **System > Services** and toggle **SSH** on.

The host system needs an established [SSH connection](#) to the remote for the rsync task. To create the connection, go to **Credentials > Backup Credentials > SSH Connections** and click **Add**. Populate the **SSH Connections** configuration panel fields as follows: Select **Semi-automatic** for the **Setup Method** and set **Private Key** to **Generate New**.

Can this be set up in a command line instead?

To use a command line, go to the **Shell** on the host system. a, enter `su - {USERNAME}`, where `{USERNAME}` is the TrueNAS user account that runs the rsync task. Enter `ssh-keygen -t rsa` to create the key pair. When prompted for a password, press Enter without setting a password (a password breaks the automated task). Here is an example of running the command:

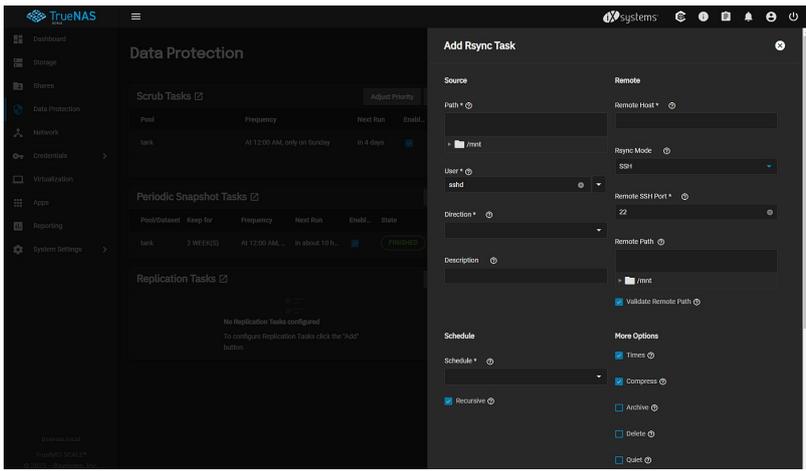
```
truenas# ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification is saved in /root/.ssh/id_rsa.
Your public key is saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:NZMgbuPvTheEqi3SA/U5wW8un6AWrx8ZsRQdbJJHmR4 tester@truenas.local
The key's randomart image is:
+---[RSA 2048]----+
|
| 0=0+
| .ooE.
| +.O=
| O.OO+. +
| ..S+. .
| ..++O.
|  oB+. .
| .=B+.o
| O+=OO
+---[SHA256]-----+
```

The default public key location is `~/ssh/id_rsa.pub`. Enter `cat ~/.ssh/id_rsa.pub` to see the key and copy the file contents. Copy it to the corresponding user account on the remote system in **Credentials > Users**. By default, SCALE only displays the root user and prompts you to display hidden users. Follow the directions to locate the **sshd** user account. Click on the **sshd** user and then on **Edit**. Paste the key in **SSH Public Key**.

Next, copy the host key from the remote system to the host system user's `~/ssh/known_hosts` directory, using `ssh-keyscan`. On the host system, open the **Shell** and enter `ssh-keyscan -t rsa {remoteIPaddress} >> {userknown_hostsDir}` where `{remoteIPaddress}` is the remote system IP address and `{userknown_hostsDir}` is the `known_hosts` directory on the host system. Example: `ssh-keyscan -t rsa 192.168.2.6 >> /root/.ssh/known_hosts`.

SSH Mode Process

Go to **Data Protection > Rsync Tasks** and click **Add**.

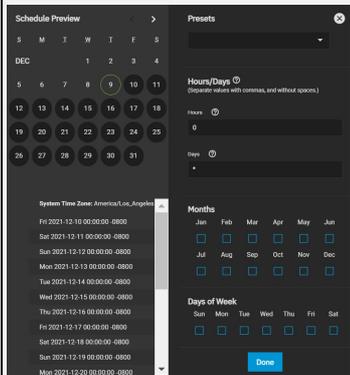


Configure the rsync task SSH settings by selecting **SSH** on the **Rsync Mode** drop-down. Type the **Port** number and **Remote Path**.

Next, define the **Source** dataset to use for the rsync task and select a **User** account. The **User** field entry must be identical to the [SSH Connection Username](#).

Choose a **Direction** for the rsync task as either **Push** or **Pull** and then define the task **Schedule**. If you need a custom schedule, select **Custom**.

Advanced Scheduler



Choosing a **Preset** option populates the rest of the fields. To customize a schedule, enter [cron](#) values for the Minutes/Hours/Days.

These fields accept standard [cron](#) values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering `10` means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means match all values.

You can set specific time ranges by entering hyphenated number values. For example, entering `30-35` in the **Minutes** field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

You can also enter lists of values. Enter individual values separated by a comma (,). For example, entering `1,14` in the **Hours** field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, entering * in **Days** runs the task every day of the month. Entering */2 runs it every other day.

Combining the above examples creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

TrueNAS has an option to select which **Months** the task runs. Leaving each month unset is the same as selecting every month.

The **Days of Week** schedules the task to run on specific days in addition to any listed days. For example, entering 1 in **Days** and setting *Wed* for **Days of Week** creates a schedule that starts a task on the first day of the month *and* every Wednesday of the month.

The **Schedule Preview** displays when the current settings mean the task runs.

Examples of CRON syntax

Syntax	Meaning	Examples
*	Every item.	* (minutes) = every minute of the hour. * (days) = every day.
*/N	Every N th item.	*/15 (minutes) = every 15th minute of the hour. */3 (days) = every 3rd day. */3 (months) = every 3rd month.
Comma and hyphen/dash	Each stated item (comma) Each item in a range (hyphen/dash).	1,31 (minutes) = on the 1st and 31st minute of the hour. 1-3,31 (minutes) = on the 1st to 3rd minutes inclusive, and the 31st minute, of the hour. mon-fri (days) = every Monday to Friday inclusive (every weekday). mar,jun,sep,dec (months) = every March, June, September, December.

You can specify days of the month or days of the week.

TrueNAS lets users create flexible schedules using the available options. The table below has some examples:

Desired schedule	Values to enter
3 times a day (at midnight, 08:00 and 16:00)	months=*; days=*; hours=0/8 or 0,8,16; minutes=0 (Meaning: every day of every month, when hours=0/8/16 and minutes=0)
Every Monday/Wednesday/Friday, at 8.30 pm	months=*; days=mon,wed,fri; hours=20; minutes=30
1st and 15th day of the month, during October to June, at 00:01 am	months=oct-dec,jan-jun; days=1,15; hours=0; minutes=1
Every 15 minutes during the working week, which is 8am - 7pm (08:00 - 19:00) Monday to Friday	Note that this requires two tasks to achieve: (1) months=*; days=mon-fri; hours=8-18; minutes=*/15 (2) months=*; days=mon-fri; hours=19; minutes=0 We need the second scheduled item, to execute at 19:00, otherwise we would stop at 18:45. Another workaround would be to stop at 18:45 or 19:45 rather than 19:00.

Next, enter the **Remote Host** IP address or hostname. Use the format *username@remote_host* if the username differs on the remote host. Configure the remaining options according to your specific needs.

Options

Source

Name	Description
Path	Browse to the path to be copied. FreeBSD file path limits apply. Other operating systems can have different limits which might affect how they can be used as sources or destinations.
User	Select the user to run the rsync task. The user selected must have permissions to write to the specified directory on the remote host.
Direction	Direct the flow of data to the remote host. During a <i>push</i> , the dataset transfers to the remote module. During a <i>pull</i> , the dataset stores files from the <i>remote</i> system.
Description	Enter a description of the rsync task.

Schedule

Name	Description
Schedule	Select a schedule preset or choose Custom to open the advanced scheduler.
Recursive	Set to include all subdirectories of the specified directory. When unset, only the specified directory is included.

Remote

Name	Description
Remote Host	Enter the IP address or hostname of the remote system that will store the copy. Use the format <i>username@remote_host</i> if the username differs on the remote host.
Rsync Mode	Choose to either use a custom-defined remote module of the rsync server or to use an SSH configuration for the rsync task.

More Options

Name	Description
Times	Set to preserve modification times of files.
Compress	Set to reduce the size of data to transmit. Recommended for slow connections.
Archive	When set, rsync is run recursively, preserving symlinks, permissions, modification times, group, and special files. When run as root, owner, device files, and special files are also preserved. Equivalent to passing the flags <code>-rptgo</code> to rsync.
Delete	Delete files in the destination directory that do not exist in the source directory.
Quiet	Set to suppress informational messages from the remote server.
Preserve Permissions	Set to preserve original file permissions. This is useful when the user is set to root.
Preserve Extended Attributes	Extended attributes are preserved, but must be supported by both systems.
Delay Updates	Set to save the temporary file from each updated file to a holding directory until the end of the transfer when all transferred files are renamed into place.
Auxiliary Parameters	Additional rsync(1) options to include. Separate entries by pressing Enter. Note: The <code>"</code> character must be escaped with a backslash (<code>\"</code> , <code>txt</code>) or used inside single quotes (<code>'*txt'</code>).
Enabled	Enable this rsync task. Unset to disable this rsync task without deleting it.

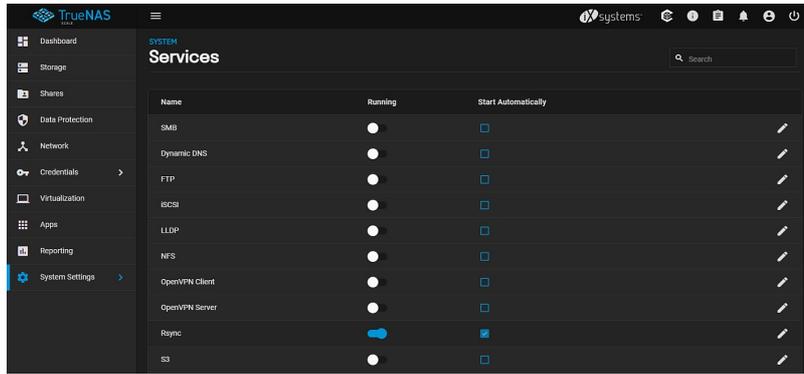
Additional options for the **SSH Rsync Mode**:

- Remote SSH Port** : Enter the SSH port number of the remote system. By default, 22 is reserved in TrueNAS.
- Remote Path** : Browse to the existing path on the remote host to sync with. Maximum path length is 255 characters.
- Validate Remote Path** : Select checkbox to automatically create the defined **Remote Path** when it does not exist.

If the **Enabled** checkbox is not selected it disables the task schedule without deleting the configuration. You can still run the rsync task by going to **Data Protection > Rsync Tasks** and clicking **Run Now** icon.

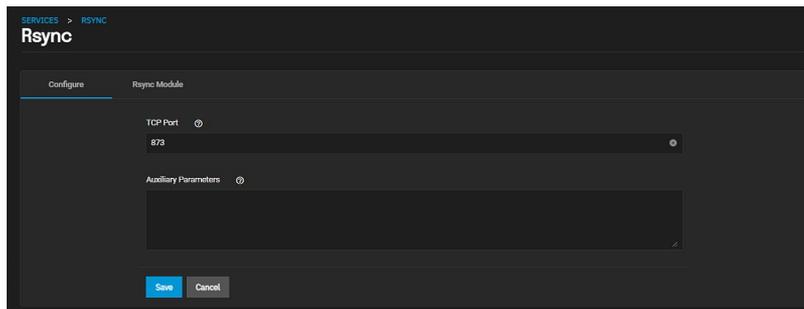
Rsync Service and Modules

The rsync task does not work unless the related system service is turned on. To turn the rsync service on, go to **System > Services** and toggle the **Rsync** on. To activate the service whenever TrueNAS boots, select the **Start Automatically** checkbox.



Click the **Configure** button to configure the service on the **Services > RSYNC > Rsync** screen. There are two tabs for rsync configuration: basic **Configure** options and **Rsync Module** creation and management.

Configure

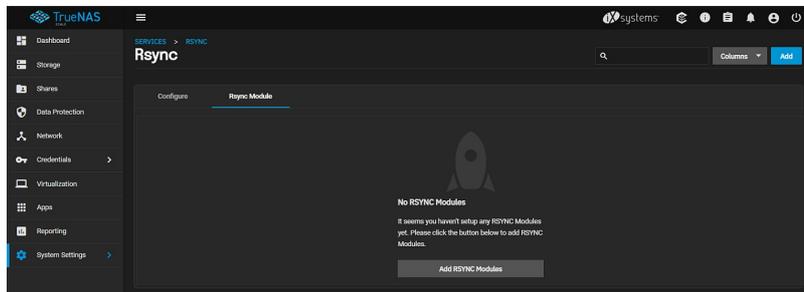


Name	Description
TCP Port	rsyncd listens on this port.
Auxiliary Parameters	Enter any additional parameters from rsyncd.conf(5) .

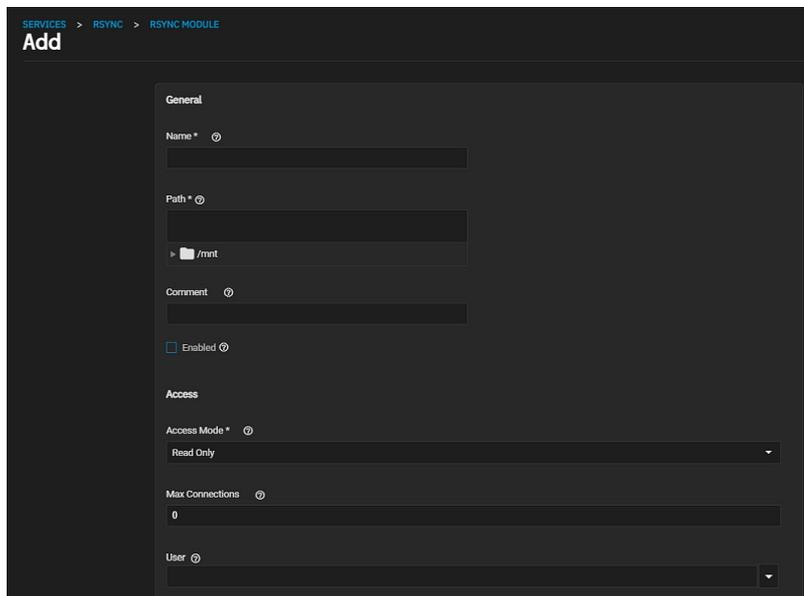
Use the default settings unless a specific change is required. Remember to click **Save** after changing any settings.

Rsync Module

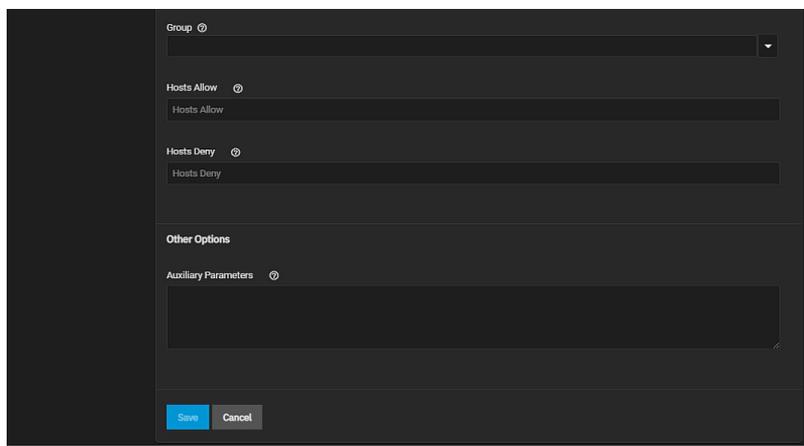
All created modules are listed on the **Rsync Module** tab.



If an rsync module is not listed, click **Add RSYNC Modules** or **Add** at the top right of the screen to add one. The **SERVICE > RSYNC > RSYNC MODULE Add** screen displays.



Scroll down to see all configuration fields.



General

Name	Description
Name	Module name that matches the name requested by the rsync client.
Path	Browse to the pool or dataset to store received data.
Comment	Describe this module.
Enabled	Activate this module for use with Rsync. Unset this field to deactivate the module without completely removing it.

Access

Name	Description
Access Mode	Choose permissions for this rsync module.
Max Connections	Maximum number of connections to this module. 0 is unlimited.
User	TrueNAS user account that runs the rsync command during file transfers to and from this module.
Group	TrueNAS group account that runs the rsync command during file transfers to and from this module.
Hosts Allow	From rsyncd.conf(5) . A list of patterns to match with the hostname and IP address of a connecting client. The connection is rejected if no patterns match. Separate entries by pressing <code>Enter</code> .
Hosts Deny	From rsyncd.conf(5) . A list of patterns to match with the hostname and IP address of a connecting client. The connection is rejected when the patterns match. Separate entries by pressing <code>Enter</code> .

Other Options

Name	Description
Auxiliary Parameters	Enter any additional parameters from rsyncd.conf(5) .

When a **Hosts Allow** list is defined, *only* the IPs and hostnames on the list are able to connect to the module.

To **EDIT** or **DELETE** a module, go to the **Rsync Modules** list and click **▶** for an entry.

6.4 - Periodic Snapshot Tasks

- [Creating a Periodic Snapshot Task](#)
 - [Process](#)

A periodic snapshot task allows scheduling the creation of read only versions of pools and datasets at a given point in time.

How should I use snapshots?

Snapshots do not make not copies of the data so creating one is quick and if little data changed, they take very little space. It is common to take frequent snapshots as soon as every 15 minutes, even for large and active pools. A snapshot where no files changed takes no storage space, but as files changes happen, the snapshot size changes to reflect the size of the changes. In the same way as all pool data, after deleting the last reference to the data you recover the space.

Snapshots keep a history of files, providing a way to recover an older copy or even a deleted file. For this reason, many administrators take snapshots often, store them for a period of time, and store them on another system, typically using the **Replication Tasks** function. Such a strategy allows the administrator to roll the system back to a specific point in time. If there is a catastrophic loss, an off-site snapshot can restore data up to the time of the last snapshot.

Creating a Periodic Snapshot Task

Any required datasets or zvols must exist before creating a snapshot task.

Process

Video Tutorial

This short video demonstrates adding a periodic snapshot task



Video Player is loading
Video URL: <https://www.tlaxnas.com/docs/files/scaleangelfishperiodicsnapshottasks.mp4>

Play Video

Play

Mute

Current Time 0:00

/

Duration 0:29

Loaded: 100.00%

Stream Type LIVE

Seek to live, currently behind liveLIVE

Remaining Time -0:29

1x

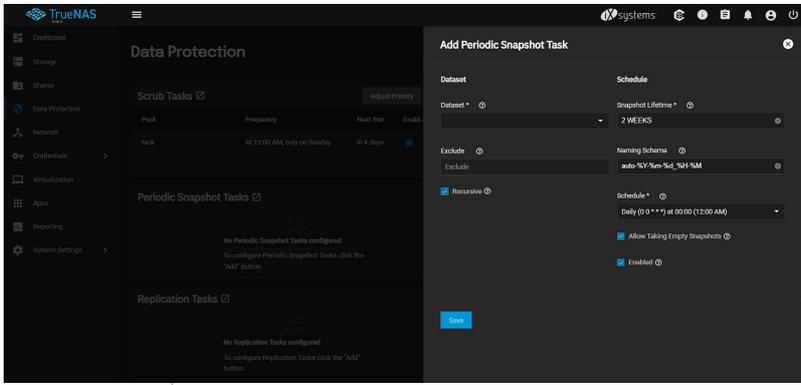
Playback Rate

Chapters

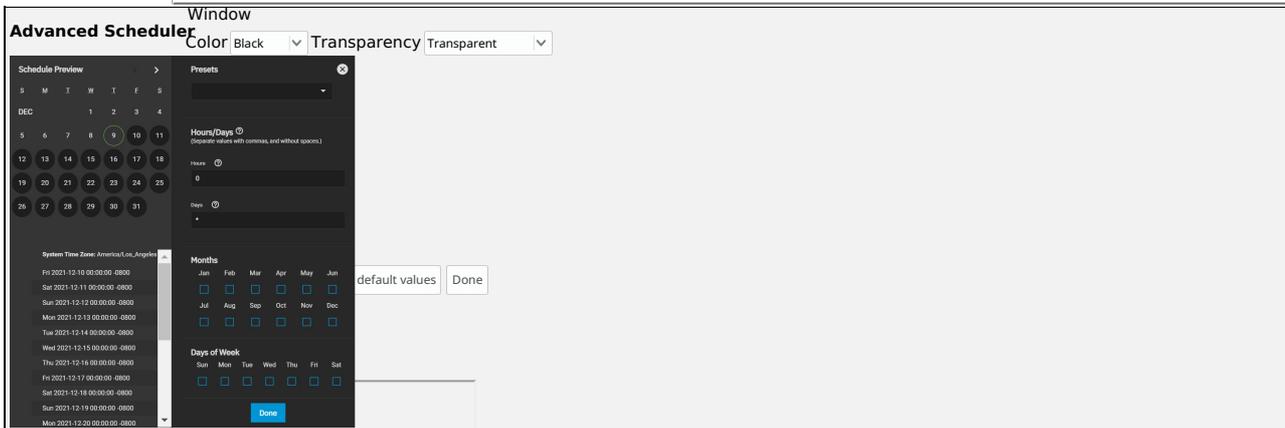
- Chapters

Descriptions

Go to **Data Protection** > **Periodic Snapshot Tasks** and click **Add**.



Choose the dataset (or zvol) to schedule as a regular back up with snapshots and how long to store snapshots. Define the task **Schedule**. If you need a specific schedule, choose **Custom** and use the **Advanced Scheduler** section below.



Choosing a **Presets** option populates the rest of the fields. To customize a schedule, enter **crontab** values for the **Minutes/Hours/Days**.

These fields accept standard **cron** values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering **10** means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means match all values.

You can set specific time ranges by entering hyphenated number values. For example, entering **30-35** in the **Minutes** field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

You can also enter lists of values. Enter individual values separated by a comma (,). For example, entering **1,14** in the **Hours** field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, entering ***** in **Days** runs the task every day of the month. Entering ***/2** runs it every other day.

Combining the above examples creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

TrueNAS has an option to select which **Months** the task runs. Leaving each month unset is the same as selecting every month.

The **Days of Week** schedules the task to run on specific days in addition to any listed days. For example, entering **1** in **Days** and setting **Wed** for **Days of Week** creates a schedule that starts a task on the first day of the month *and* every Wednesday of the month.

The **Schedule Preview** displays when the current settings mean the task runs.

Examples of CRON syntax

Syntax	Meaning	Examples
*	Every item.	* (minutes) = every minute of the hour. * (days) = every day.
*/N	Every N th item.	*/15 (minutes) = every 15th minute of the hour. */3 (days) = every 3rd day. */3 (months) = every 3rd month.
Comma and hyphen/dash	Each stated item (comma) Each item in a range (hyphen/dash).	1,31 (minutes) = on the 1st and 31st minute of the hour. 1-3,31 (minutes) = on the 1st to 3rd minutes inclusive, and the 31st minute, of the hour. mon-fri (days) = every Monday to Friday inclusive (every weekday). mar,jun,sep,dec (months) = every March, June, September, December.

You can specify days of the month or days of the week.

TrueNAS lets users create flexible schedules using the available options. The table below has some examples:

Desired schedule	Values to enter
3 times a day (at midnight, 08:00 and 16:00)	months=*; days=*; hours=0/8 or 0,8,16; minutes=0 (Meaning: every day of every month, when hours=0/8/16 and minutes=0)
Every Monday/Wednesday/Friday, at 8.30 pm	months=*; days=mon,wed,fri; hours=20; minutes=30
1st and 15th day of the month, during October to June, at 00:01 am	months=oct-dec,jan-jun; days=1,15; hours=0; minutes=1
Every 15 minutes during the working week, which is 8am - 7pm (08:00 - 19:00) Monday to Friday	Note that this requires two tasks to achieve: (1) months=*; days=mon-fri; hours=8-18; minutes=*/15 (2) months=*; days=mon-fri; hours=19; minutes=0 We need the second scheduled item, to execute at 19:00, otherwise we would stop at 18:45. Another workaround would be to stop at 18:45 or 19:45 rather than 19:00.

Configure the remaining options for your use case.

Specific Options	
Dataset	
Name	Description
Dataset	Select a pool, dataset, or zvol.
Recursive	Set to take separate snapshots of the dataset and each of its child datasets. Leave unset to take a single snapshot only of the specified dataset

	without child datasets.
Exclude	Exclude specific child datasets from the snapshot. Use with recursive snapshots. List paths to any child datasets to exclude. Example: pool1/dataset1/child1. A recursive snapshot of pool1/dataset1 will include all child datasets except child1. Separate entries by pressing Enter.

Schedule

Name	Description
Snapshot Lifetime	Define a length of time to retain the snapshot on this system using a numeric value and a single lowercase letter for units. Examples: <i>3h</i> is three hours, <i>1m</i> is one month, and <i>1y</i> is one year. Does not accept Minute values. After the time expires, the snapshot is removed. Snapshots which have been replicated to other systems are not affected.
Naming Schema	Snapshot name format string. The default is <code>auto-%Y-%m-%d-%H-%M</code> . Must include the strings <code>%Y</code> , <code>%m</code> , <code>%d</code> , <code>%H</code> , and <code>%M</code> , which are replaced with the four-digit year, month, day of month, hour, and minute as defined in strftime(3) . For example, snapshots of <i>pool1</i> with a Naming Schema of <code>customsnap-%Y%m%d.%H%M</code> have names like <code>pool1@customsnap-20190315.0527</code> .
Schedule	Choose one of the presets or <i>Custom</i> to use the advanced scheduler.
Allow Taking Empty Snapshots	Creates dataset snapshots even when there have been no changes to the dataset from the last snapshot. Recommended for long-term restore points, multiple snapshot tasks pointed at the same datasets, or compatibility with snapshot schedules or replications created in TrueNAS 11.2 and earlier. For example, allowing empty snapshots for a monthly snapshot schedule allows that monthly snapshot to be taken, even when a daily snapshot task has already taken a snapshot of any changes to the dataset.
Enabled	To activate this periodic snapshot schedule, set this option. To disable this task without deleting it, unset this option.

Naming Schemas

The **Naming Schema** determines how automated snapshot names generate. A valid schema requires the `%Y` (year), `%m` (month), `%d` (day), `%H` (hour), and `%M` (minute) time strings, but you can add more identifiers to the schema too, using any identifiers from the Python [strftime function](#).

For **Periodic Snapshot Tasks** used to set up a replication task with the **Replication Task** function:

You can use custom naming schema for full backup replication tasks. If you are going to use the snapshot for an incremental replication task, use the default naming schema. Go to [Using a Custom Schema](#) for additional information.

This uses some letters differently from POSIX (Unix) time functions. For example, including `%z` (time zone) ensures that snapshots do not have naming conflicts when daylight time starts and ends, and `%S` (second) adds finer time granularity.

Examples:

Naming Scheme	Snapshot Names Look Like
<code>replicationsnaps-1wklife-%Y%m%d_%H:%M</code>	<code>replicationsnaps-1wklife-20210120_00:00</code> , <code>replicationsnaps-1wklife-20210120_06:00</code>
<code>autosnap_%Y.%m.%d-%H.%M.%S-%z</code>	<code>autosnap_2021.01.20-00.00-EST</code> , <code>autosnap_2021.01.20-06.00-EST</code>

When referencing snapshots from a Windows computer, avoid using characters like `:` that are invalid in a Windows file path. Some applications limit filename or path length, and there might be limitations related to spaces and other characters. Always consider future uses and ensure the name given to a periodic snapshot is acceptable.

Snapshot Lifetimes

TrueNAS deletes snapshots when they reach the end of their life and preserves snapshots when at least one periodic task requires it. For example, you have two schedules created where one schedule takes a snapshot every hour and keeps them for a week, and the other takes a snapshot every day and keeps them for 3 years. Each has an hourly snapshot taken. After a week, snapshots created at `01.00` through `23.00` get deleted, but you keep snapshots timed at `00.00` because they are necessary for the second periodic task. These snapshots get destroyed at the end of 3 years.

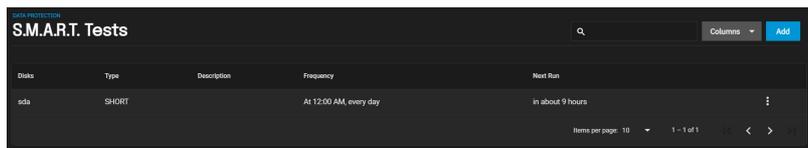
Click **Save** to save this task and add it to the list in **Data Protection > Periodic Snapshot Tasks**. You can find any snapshots taken using this task in **Storage > Snapshots**.

To check the log for a saved snapshot schedule, go to **Data Protection > Periodic Snapshot Tasks** and click on the task. The **Edit Periodic Snapshot Tasks** screen displays where you can modify any settings for the task.

6.5 - S.M.A.R.T. Tests

- Manual S.M.A.R.T. Test
 - Automatic S.M.A.R.T. Tests
 - Service Options

S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) is a standard for disk monitoring and testing. You can monitor disks for problems using different kinds of self-tests. TrueNAS can adjust when it issues S.M.A.R.T. alerts. When S.M.A.R.T. monitoring reports a disk issue, we recommend you replace that disk. Most modern ATA, IDE, and SCSI-3 hard drives support S.M.A.R.T. Refer to your respective drive documentation for confirmation.



TrueNAS runs S.M.A.R.T. tests on disks. Running tests can reduce drive performance, so we recommend scheduling tests when the system is in a low-usage state. Avoid scheduling disk-intensive tests at the same time! For example, don't schedule S.M.A.R.T. tests on the same day as a disk [scrub](#) or other Data Protection task.

How do I check or change S.M.A.R.T. testing for a disk?

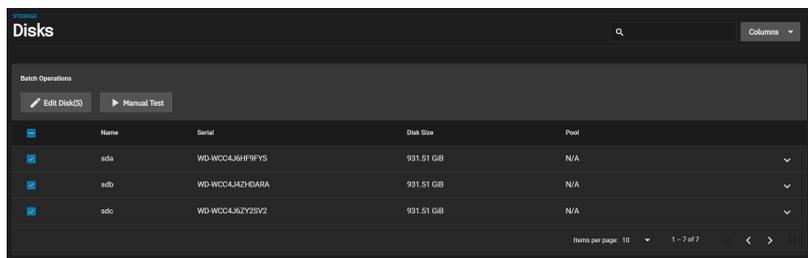
Go to **Storage**, then click the **Disks** drop-down and select **Disks**.

Click the in a disk's row to expand it. *Enable S.M.A.R.T.* shows as *true* or *false*.

To enable or disable testing, click **EDIT** and find the **Enable S.M.A.R.T.** option.

Manual S.M.A.R.T. Test

To quickly test one or more disk for errors, select the disks you want to test and click **MANUAL TEST**.



Next, select the test **Type**. Test types differ based on the drive connection, ATA or SCSI:

ATA

- Long** runs a S.M.A.R.T. Extended Self Test that scans the entire disk surface, which may take hours on large-volume disks.
- Short** runs a basic S.M.A.R.T. Short Self Test (usually under ten minutes) that varies by manufacturer.
- Conveyance** runs a S.M.A.R.T. Conveyance Self Test (usually only minutes) that identifies damage incurred while transporting the device.
- Offline** runs a S.M.A.R.T. Immediate Offline Test that updates the S.M.A.R.T. Attribute values. Errors will appear in the S.M.A.R.T. error log.

SCSI

- Long** runs the "Background long" self-test.
- Short** runs the "Background short" self-test.
- Offline** runs the default self-test in the foreground, but doesn't place an entry in the self-test log.

For more information, refer to [smartctl\(8\)](#).

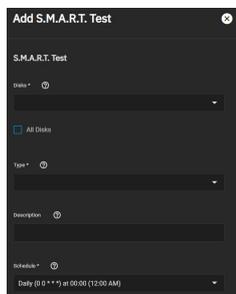
Click **START** to begin the test. Test duration varies based on the test type you chose. TrueNAS generates alerts when tests discover issues.

Where can I view the test results?

Click the in a disk's row to expand it, then click **S.M.A.R.T. TEST RESULTS**. You can also see results in the [Shell](#) using `smartctl -l selftest /dev/ada0`.

Automatic S.M.A.R.T. Tests

To schedule recurring S.M.A.R.T. tests, go to **Data Protection** and click **ADD** in the **S.M.A.R.T. Tests** window.



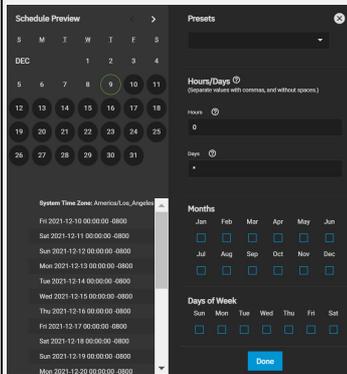
Specific Options	
Name	Description
All Disks	Setting <i>All Disks</i> includes every disk with S.M.A.R.T. enabled. Leave unset to choose which <i>Disks</i> to test.
Disks	Select the disks to monitor.
Type	Choose the test type. See smartctl(8) for descriptions of each type. Some types degrade performance or take disks offline.
Description	Enter information about the S.M.A.R.T. test.
Schedule	The time the test runs. Choose a preset or select <i>Custom</i> to open the advanced scheduler.

Choose the **Disks** to test, the test **Type** to run, and the task's **Schedule**.

S.M.A.R.T. tests can offline disks! Avoid scheduling S.M.A.R.T. tests simultaneously with scrub or other data protection tasks.

If you need the test to run on a specific **Schedule**, choose **Custom** to open the advanced scheduler.

Advanced Scheduler



Choosing a **Presets** option populates in the rest of the fields. To customize a schedule, enter **crontab** values for the Minutes/Hours/Days.

These fields accept standard **cron** values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering **10** means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means match all values.

You can set specific time ranges by entering hyphenated number values. For example, entering **30-35** in the **Minutes** field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

You can also enter lists of values. Enter individual values separated by a comma (.). For example, entering **1,14** in the **Hours** field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, entering ***** in **Days** runs the task every day of the month. Entering ***/2** runs it every other day.

Combining the above examples creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

TrueNAS has an option to select which **Months** the task runs. Leaving each month unset is the same as selecting every month.

The **Days of Week** schedules the task to run on specific days in addition to any listed days. For example, entering **1** in **Days** and setting **Wed** for **Days of Week** creates a schedule that starts a task on the first day of the month *and* every Wednesday of the month.

The **Schedule Preview** displays when the current settings mean the task runs.

Examples of CRON syntax

Syntax	Meaning	Examples
*	Every item.	* (minutes) = every minute of the hour. * (days) = every day.
*/N	Every N th item.	*/15 (minutes) = every 15th minute of the hour. */3 (days) = every 3rd day. */3 (months) = every 3rd month.
Comma and hyphen/dash	Each stated item (comma) Each item in a range (hyphen/dash).	1,31 (minutes) = on the 1st and 31st minute of the hour. 1-3,31 (minutes) = on the 1st to 3rd minutes inclusive, and the 31st minute, of the hour. mon-fri (days) = every Monday to Friday inclusive (every weekday). mar,jun,sep,dec (months) = every March, June, September, December.

You can specify days of the month or days of the week.

TrueNAS lets users create flexible schedules using the available options. The table below has some examples:

Desired schedule	Values to enter
3 times a day (at midnight, 08:00 and 16:00)	months=*; days=*; hours=0/8 or 0,8,16; minutes=0 (Meaning: every day of every month, when hours=0/8/16 and minutes=0)
Every Monday/Wednesday/Friday, at 8.30 pm	months=*; days=mon,wed,fri; hours=20; minutes=30
1st and 15th day of the month, during October to June, at 00:01 am	months=oct-dec,jan-jun; days=1,15; hours=0; minutes=1
Every 15 minutes during the working week, which is 8am - 7pm (08:00 - 19:00) Monday to Friday	Note that this requires two tasks to achieve: (1) months=*; days=mon-fri; hours=8-18; minutes=*/15 (2) months=*; days=mon-fri; hours=19; minutes=0 We need the second scheduled item, to execute at 19:00, otherwise we would stop at 18:45. Another workaround would be to stop at 18:45 or 19:45 rather than 19:00.

Saved schedules appear in the **S.M.A.R.T. Tests** window.

CLI

To verify the schedule is saved, you can open the [shell](#) and enter `smartd -q showtests`.

Service Options

You must enable the S.M.A.R.T. service to run automatic S.M.A.R.T. tests.

RAID controllers?

Disable the S.M.A.R.T. service when a RAID controller controls the disks. The controller monitors S.M.A.R.T. separately and marks disks as a **Predictive Failure** on a test failure.

To start the S.M.A.R.T. service, go to **System Settings > Services** and toggle **S.M.A.R.T.**. To start the service during the TrueNAS boot process, set **Start Automatically**.

Configure the S.M.A.R.T. service by clicking

General Options

Check Interval * 30

Power Mode * Never

Difference * 0

Informational * 0

Critical * 0

Save

Name	Description
Check Interval	Minutes for smartd to wake up and check if any tests should run.
Power Mode	S.M.A.R.T. only tests when the Power Mode is Never.
Difference	Degrees in Celsius. S.M.A.R.T. reports if a drive's temperature has changed by N degrees Celsius since the last report.
Informational	Threshold temperature in Celsius. S.M.A.R.T. will message with a LOG_INFO log level if the temperature is above the threshold.
Critical	Threshold temperature in Celsius. S.M.A.R.T. will message with a LOG_CRIT log level and send an email if the temperature is above the threshold.

Click **Save** after changing any settings.

6.6 - Replication

Remote Replication

- - [Remote Replication](#)
 - [Creating a Remote Replication Task](#)
 - [Local Replication](#)
 - [Quick Backups with the Replication Wizard](#)
 - [Advanced Replication](#)
 - [Creating an Advanced Replication Task](#)
 - [Unlocking a Replicated Encrypted Dataset or Zvol Without a Passphrase](#)

Configure SSH in TrueNAS before creating a remote replication task. This ensures that new snapshots are regularly available for replication.

To streamline creating simple replication configurations, the replication wizard assists with creating a new SSH connection and automatically creates a periodic snapshot task for sources that have no existing snapshots.

Process Summary

Process Summary

- **Data Protection > Replication Tasks**
 - Choose sources for snapshot replication.
 - Remote sources require an SSH connection.
 - TrueNAS shows the number snapshots available to replicate.
- Define the snapshot destination.
 - A remote destination requires an SSH connection.
 - Choose destination or define manually by typing a path.
 - Adding a new name on the end of the path creates a new dataset.
- Choose replication security.
 - iXsystems always recommend replication with encryption.
 - Disabling encryption is only meant for absolutely secure and trusted destinations.
- Schedule the replication.
 - You can schedule standardized presets or a custom defined schedule.
 - Running once runs the replication immediately after creation.
 - Task is still saved and you can rerun or edit it.
- Choose how long to keep the replicated snapshots.

Replication Tutorial

This video tutorial presents a simple example of setting up replication.



Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaleangelfishreplication.mp4>

Play Video

Play

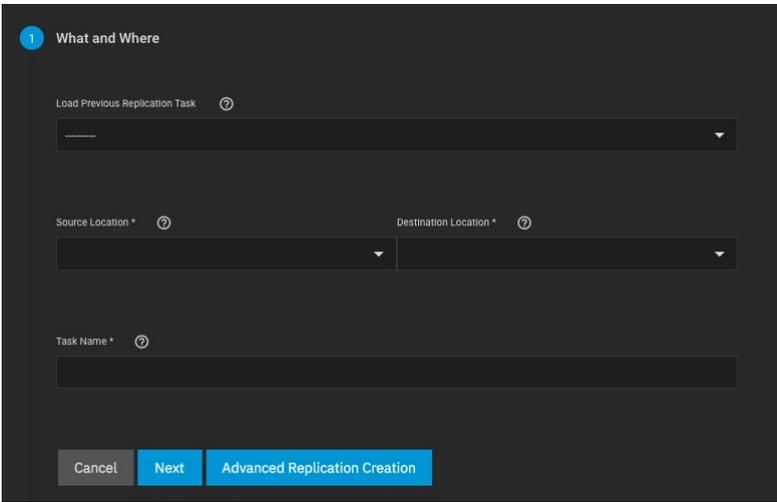
Creating a Remote Replication Task

To create a new replication, go to **Data Protection > Replication Tasks** and click **ADD**.

Current Time 0:00

Duration 1:04

Loaded: 100.00%



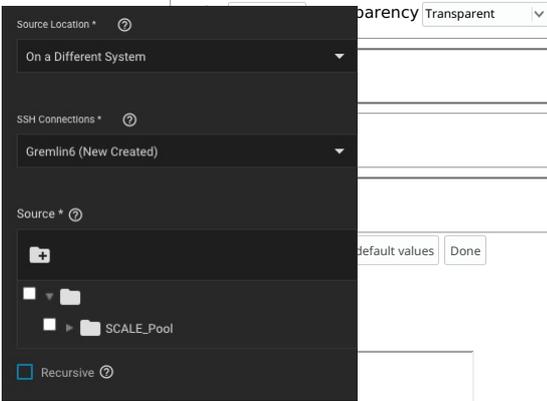
You can load any saved replication to prepopulate the wizard with that configuration. Saving changes to the configuration creates a new replication task without altering the task you loaded into the wizard. This saves some time when creating multiple replication tasks between the same two systems.

Sources

Beginning of dialog window. Escape will cancel and close the window.

Start by configuring the application sources. Sources are the datasets or zvols with snapshots to use for replication. Choosing a remote source requires selecting an SSH connection to that system. Expanding the directory browser shows the current datasets or zvols that are available for replication. You can select multiple sources or manually type a path in the field.

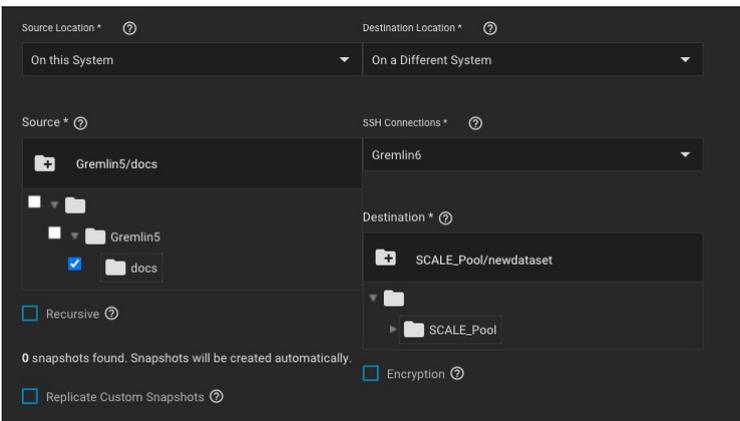
TrueNAS shows how many snapshots are available for replication. We recommend you manually snapshot the sources or create a periodic snapshot task before creating the replication task. However, when the sources are on the local system and don't have any existing snapshots, TrueNAS can create a basic periodic snapshot task and snapshot the sources immediately before starting the replication. Enabling **Recursive** replicates all snapshots contained within the selected source dataset snapshots.



Local sources can also use a naming schema to identify any custom snapshots to include in the replication. Remote sources require entering a *snapshot naming schema* to identify the snapshots to replicate. A naming schema is a collection of [strftime](#) time and date strings and any identifiers that a user might have added to the snapshot name.

Destination

The destination is where replicated snapshots are stored. Choosing a remote destination requires an SSH connection to that system. Expanding the directory browser shows the current datasets that are available for replication. You can select a destination dataset or manually type a path in the field. You cannot use zvols as a remote replication destination. Adding a name to the end of the path creates a new dataset in that location.



To use encryption when replicating data click the **Encryption** box. After selecting the box these additional encryption options become available:

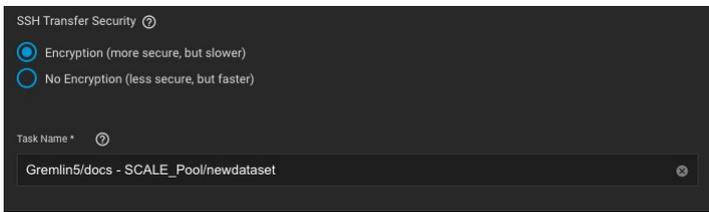
- **Encryption Key Format** allows the user to choose between a hex (base 16 numeral) or passphrase (alphanumeric) style encryption key.
- **Store Encryption key in Sending TrueNAS database** allows the user to either store the encryption key in the sending TrueNAS database (box checked) or choose a temporary location for the encryption key that decrypts replicated data (box unchecked)

Security and Task Name

Using encryption for SSH transfer security is always recommended.

In situations where two systems within an absolutely secure network are used for replication, disabling encryption speeds up the transfer. However, the data is completely unprotected from eavesdropping.

Choosing **no encryption** for the task is less secure but faster. This method uses common port settings but these can be overridden by switching to the advanced options screen or editing the task after creation.

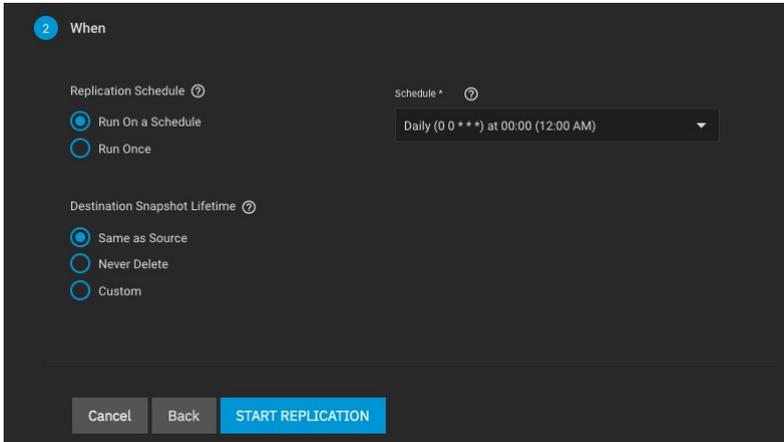


TrueNAS suggests a name based off the selected sources and destination, but this can be overwritten with a custom name.

Schedule and Lifetime

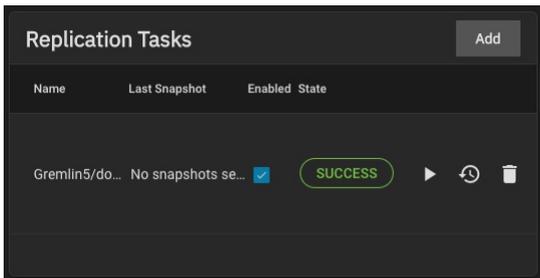
Adding a schedule automates the task to run according to your chosen times. You can choose between a number of preset schedules or create a custom schedule for when the replication runs. Choosing to run the replication once runs the replication immediately after saving the task, but you must manually trigger any additional replications.

Finally, define how long you want to keep snapshots on the destination system. We generally recommend defining snapshot lifetime to prevent cluttering the system with obsolete snapshots.

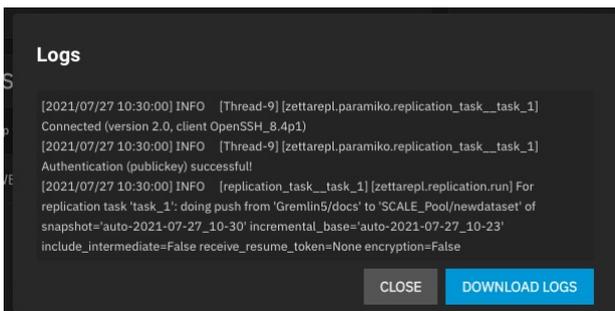


Starting the Replication

Start Replication* saves the new replication task. New tasks are enabled by default and activate according to their schedule or immediately when no schedule is chosen. The first time a replication task runs, it takes longer because the snapshots must be copied entirely fresh to the destination.



Later replications run faster, as only the subsequent changes to snapshots are replicated. Clicking the task state opens the log for that task.



Local Replication

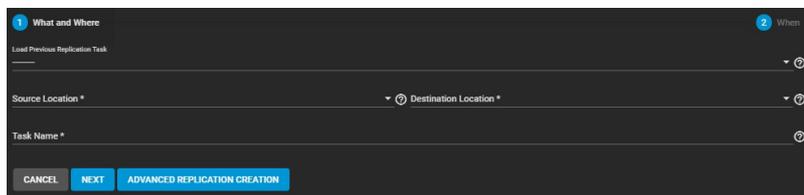
Process Summary

Process Summary

- Requirements: Storage pools and datasets created in **Storage > Pools**.
- Go to **Data Protection > Replication Tasks** and click **ADD**
 - Choose **Sources**
 - Set the source location to the local system
 - Use the file browser or type paths to the sources
 - Define a **Destination** path
 - Set the destination location to the local system
 - Select or manually define a path to the single destination location for the snapshot copies.
 - Set the **Replication schedule** to run once
 - Define how long the snapshots are stored in the **Destination**
 - Clicking **START REPLICATION** immediately snapshots the chosen sources and copies those snapshots to the destination
 - Dialog might ask to delete existing snapshots from the destination. Be sure that all important important data is protected before deleting anything.
- Clicking the task **State** shows the logs for that replication task.

Quick Backups with the Replication Wizard

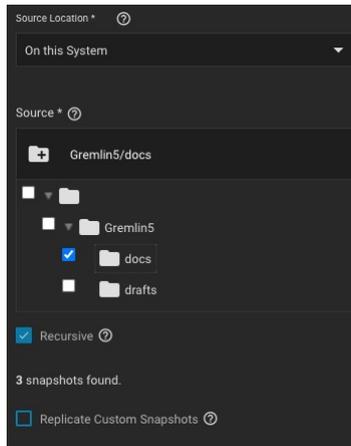
TrueNAS provides a wizard for quickly configuring different simple replication scenarios.



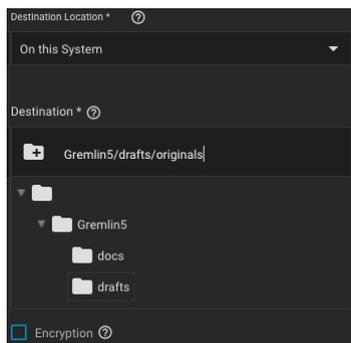
While we recommend regularly scheduled replications to a remote location as the optimal backup scenario, the wizard can very quickly create and copy ZFS snapshots to another location on the same system. This is useful when no remote backup locations are available, or when a disk is in immediate danger of failure.

The only thing you need before creating a quick local replication are datasets or zvols in a storage pool to use as the replication source and (preferably) a second storage pool to use for storing replicated snapshots. You can set up the local replication entirely in the **Replication Wizard**.

To open the **Replication Wizard**, go to **Data Protection > Replication Tasks** and click **ADD**. Set the source location to the local system and pick which datasets to snapshot. The wizard takes new snapshots of the sources when no existing source snapshots are found. Enabling **Recursive** replicates all snapshots contained within the selected source dataset snapshots. Local sources can also use a naming schema to identify any custom snapshots to include in the replication. A naming schema is a collection of [strftime](#) time and date strings and any identifiers that a user might have added to the snapshot name.



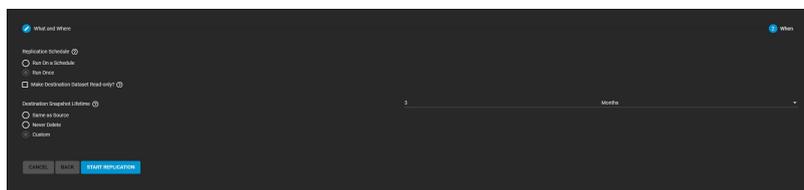
Set the destination to the local system and define the path to the storage location for replicated snapshots. When manually defining the destination, be sure to type the full path to the destination location.



TrueNAS suggests a default name for the task based on the selected source and destination locations, but you can type your own name for the replication. You can load any saved replication task into the wizard to make creating new replication schedules even easier.

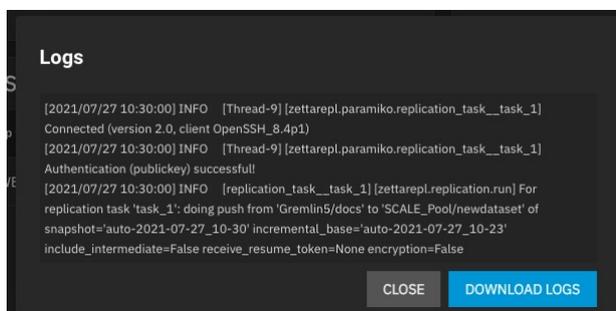
You can define a specific schedule for this replication or choose to run it immediately after saving the new task. Unscheduled tasks are saved in the replication task list and you can run saved tasks manually or edit later to add a schedule.

The destination lifetime is how long copied snapshots are stored in the destination before they are deleted. We usually recommend defining a snapshot lifetime to prevent storage issues. Choosing to keep snapshots indefinitely can require you to manually clean old snapshots from the system if or when the destination fills to capacity.



Clicking **START REPLICATION** saves the new task and immediately attempts to replicate snapshots to the destination. When TrueNAS detects that the destination already has unrelated snapshots, it asks to delete the unrelated snapshots and do a full copy of the new snapshots. This can delete important data, so be sure any existing snapshots can be deleted or are backed up in another location.

The simple replication is added to the replication task list and shows that it is currently running. Clicking the task state shows the replication log with an option to download the log to your local system.



To confirm that snapshots are replicated, go to **Storage > Snapshots > Snapshots** and verify the destination dataset has new snapshots with correct timestamps.

Dataset	Snapshot	Used	Date Created	Referenced
tank	auto-2021-08-10	80.00 KIB	2021-08-10	96.00 KIB

Advanced Replication

Requirements:

- Storage pools with datasets and data to snapshot.
- SSH configured with a connection to the remote system saved in **Credentials > Backup Credentials > SSH Connections**.
- Dataset snapshot task saved in **Data Protection > Periodic Snapshot Tasks**.

Process Summary

Go to **Data Protection > Replication Tasks** and click **ADD**, then select **ADVANCED REPLICATION CREATION**.

- General Options:
 - Name the task.
 - Select **Push** or **Pull** for the local system.
 - Select a replication transport method.
 - SSH is recommended.
 - SSH+Netcat is used for secured networks.
 - Local is for in-system replication.
- Configure the replication transport method:
 - Remote options require a preconfigured SSH connection.
 - SSH+Netcat requires defining netcat ports and addresses.
- Sources:
 - Select sources for replication.
 - Choose a preconfigured periodic snapshot task as the source of snapshots to replicate.
 - Remote sources require defining a snapshot naming schema.
- Destination:
 - Remote destination requires an SSH connection.
 - Select a destination or type a path in the field.
 - Define how long to keep snapshots in the destination.
- Scheduling:
 - Run automatically starts the replication after a related periodic snapshot task completes.
 - To automate the task according to its own schedule, set the *schedule* option and define a schedule for the replication task.

Creating an Advanced Replication Task

To use the advanced editor to create a replication task, go to **Data Protection > Replication Tasks**, click **ADD** to open the wizard, then click the **ADVANCED REPLICATION CREATION** button.

Options are grouped together by category. Options can appear, disappear, or be disabled depending on the configuration choices you make. Start by configuring the **General** options first, then the **Transport** options before configuring replication **Source**, **Destination**, and **Replication Schedule**.

General

Name *

Direction

Transport

Number of retries for failed replications

Logging Level

Enabled

Type a name for the task in **Name**. Each task name must be unique, and we recommend you name it in a way that makes it easy to remember what the task is doing.

Direction allows you to choose whether the local system is sending (**Push**) or receiving data (**Pull**).

Decide what **Transport** method (SSH, SSH+NETCAT, or LOCAL) to use for the replication before configuring the other sections.

Set the **Number of retries for failed replications** before stopping and marking the task as failed (the default is 5).

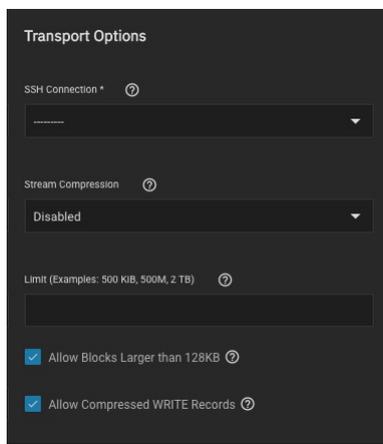
Use the **Logging Level** to set the message verbosity level in the replication task log.

To ensure the replication task is active check the **Enabled** box.

Transport Options

The **Transport** selector determines the method to use for the replication: **SSH** is the standard option for sending or receiving data from a remote system, but **SSH+NETCAT** is available as a faster option for replications that take place within completely secure networks. **Local** is only used for replicating data to another location on the same system.

With SSH-based replications, configure the transport method by selecting the **SSH Connection** to the remote system that sends or receives snapshots. Options for compressing data, adding a bandwidth limit, or other data stream customizations are available. **Stream Compression** options are only available when using SSH. Before enabling **Compressed WRITE Records**, verify that the destination system also supports compressed WRITE records.



For SSH+NETCAT replications, you must define the addresses and ports to use for the Netcat connection.

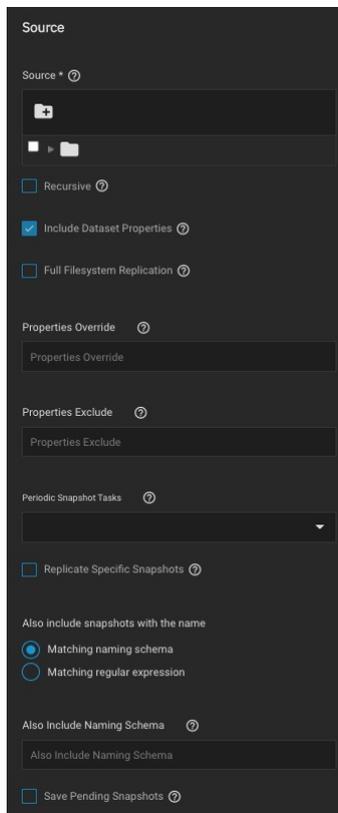
Allow Blocks Larger than 128KB is a one-way toggle. Replication tasks using large block replication only continues to work as long as this option remains enabled.

Source

The replication **Source** is the datasets or zvols to use for replication. Select the sources to use for this replication task by opening the file browser or entering dataset names in the field. Pulling snapshots from a remote source requires a valid **SSH Connection** before the file browser can show any directories.

If the file browser shows a connection error after selecting the correct **SSH Connection**, you might need to log in to the remote system and make sure it is configured to allow SSH connections.

In TrueNAS, do this by going to the **System Settings > Services** screen, checking the **SSH** service configuration, and starting the service.



By default, the replication task uses snapshots to quickly transfer data to the receiving system. When **Full Filesystem Replication** is set, the chosen **Source** is completely replicated, including all dataset properties, snapshots, child datasets, and clones. When choosing this option, it is recommended to allocate additional time for the replication task to run.

Leaving **Full Filesystem Replication** unset but setting **Include Dataset Properties** includes just the dataset properties in the snapshots to be replicated.

Checking the **Recursive** check box allows you to recursively replicate child dataset snapshots or exclude specific child datasets or properties from the replication.

Enter new defined properties in the **Properties Override** field to replace existing dataset properties with the newly defined properties in the replicated files.

List any existing dataset properties to remove from the replicated files in the **Properties Exclude** field.

Local sources are replicated by snapshots that were generated from a periodic snapshot task and/or from a defined naming schema that matches manually created snapshots.

Select a previously configured periodic snapshot task for this replication task in **Periodic Snapshot Tasks** drop down list. The replication task selected must have the same vales in **Recursive** and **Exclude Child Datasets** as the chosen periodic snapshot task. Selecting a periodic snapshot schedule removes the **Schedule** field.

To define specific snapshots from the periodic task to use for the replication, set **Replicate Specific Snapshots** and enter a schedule. The only periodically generated snapshots included in the replication task are those that match your defined schedule.

Remote sources require entering a snapshot naming schema to identify the snapshots to replicate. A naming schema is a collection of [strftime](#) time and date strings and any identifiers that a user might have added to the snapshot name. For example, entering the naming schema `custom-%Y-%m-%d_%H-%M` finds and replicates snapshots like `custom-2020-03-25_09-15`. Multiple schemas can be entered by pressing Enter to separate each schema.

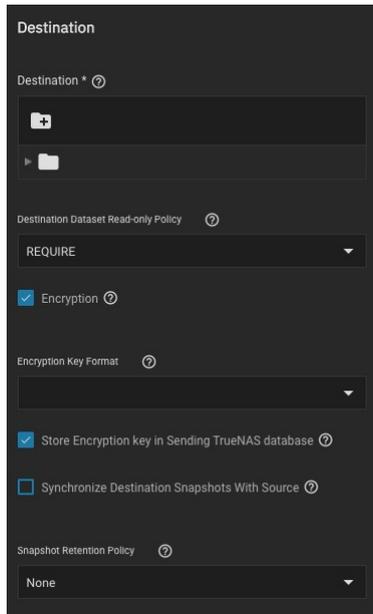
Alternately, you can use your **Replication Schedule** to determine which snapshots are replicated by setting **Run Automatically, Only Replicate Snapshots Matching Schedule**, and defining when the replication task runs.

When a replication task is having difficulty completing, it is a good idea to set **Save Pending Snapshots**. This prevents the source TrueNAS from automatically deleting any snapshots that fail to replicate to the destination system.

Destination

Use **Destination** to specify where replicated data is stored. Choosing a remote destination requires an [SSH Connection](#) to that system. Expanding the file browser shows the current datasets that are available on the destination system. You can click a destination or manually type a path in the field. Adding a name to the end of the path creates a new dataset in that location.

DO NOT use zvols for a remote destination



By default, the destination dataset is set to be read-only* after the replication is complete. You can change the **Destination Dataset Read-only Policy** to only start replication when the destination is read-only (**REQUIRE**) or to disable checking the dataset's read-only state (**IGNORE**).

The **Encryption** checkbox adds another layer of security to replicated data by encrypting the data before transfer and decrypting it on the destination system.

- Setting the checkbox adds more options to choose between using a **HEX** key or defining your own encryption **PASSPHRASE**.
- You can store the encryption key either in the TrueNAS system database or in a custom-defined location.

Synchronizing Destination Snapshots With Source destroys any snapshots in the destination that do not match the source snapshots. TrueNAS also does a full replication of the source snapshots as if the replication task had never been run before, which can lead to excessive bandwidth consumption.

This can be a very destructive option. Make sure that any snapshots deleted from the destination are obsolete or otherwise backed up in a different location.

Defining the **Snapshot Retention Policy** is generally recommended to prevent cluttering the system with obsolete snapshots. Choosing **Same as Source** keeps the snapshots on the destination system for the same amount of time as the defined **Snapshot Lifetime** from the source system periodic snapshot task.

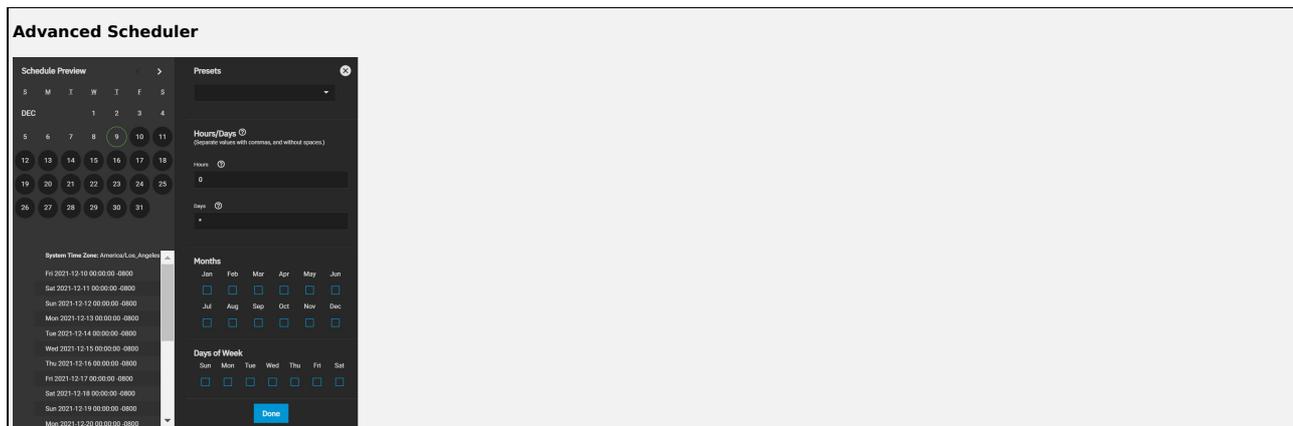
You can use **Custom** to define your own lifetime for snapshots on the destination system.

Schedule

By default, setting the task to **Run Automatically** starts the replication immediately after the related periodic snapshot task is complete.

Setting the **Schedule** checkbox allows scheduling the replication to run at a separate time.

- Defining a specific time for the replication task to run is a must do.
- Choose a time frame that both gives the replication task enough time to finish and is during a time of day when network traffic for both source and destination systems is minimal.
- Use the custom scheduler (recommended) when you need to fine-tune an exact time or day for the replication.



Choosing a **Presets** option populates in the rest of the fields. To customize a schedule, enter [cron](#) values for the Minutes/Hours/Days.

These fields accept standard [cron](#) values. The simplest option is to enter a single number in the field. The task runs when the time value matches that number. For example, entering **10** means that the job runs when the time is ten minutes past the hour.

An asterisk (*) means match all values.

You can set specific time ranges by entering hyphenated number values. For example, entering **30-35** in the **Minutes** field sets the task to run at minutes 30, 31, 32, 33, 34, and 35.

You can also enter lists of values. Enter individual values separated by a comma (,). For example, entering **1,14** in the **Hours** field means the task runs at 1:00 AM (0100) and 2:00 PM (1400).

A slash (/) designates a step value. For example, entering * in **Days** runs the task every day of the month. Entering */2 runs it every other day.

Combining the above examples creates a schedule running a task each minute from 1:30-1:35 AM and 2:30-2:35 PM every other day.

TrueNAS has an option to select which **Months** the task runs. Leaving each month unset is the same as selecting every month.

The **Days of Week** schedules the task to run on specific days in addition to any listed days. For example, entering **1** in **Days** and setting **Wed** for **Days of**

Week creates a schedule that starts a task on the first day of the month *and* every Wednesday of the month.

The **Schedule Preview** displays when the current settings mean the task runs.

Examples of CRON syntax

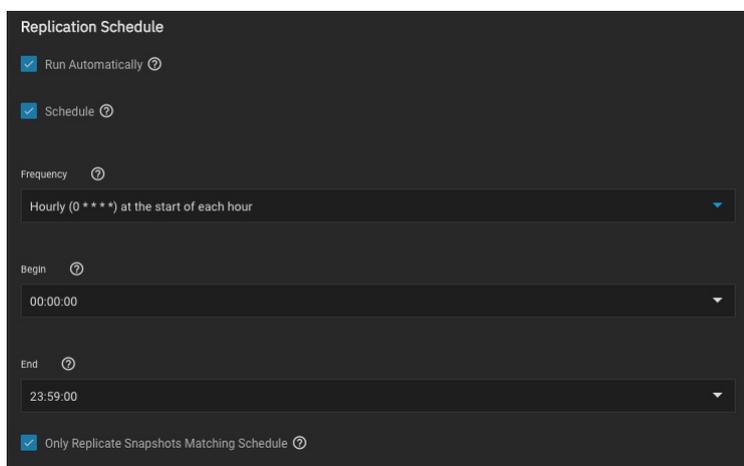
Syntax	Meaning	Examples
*	Every item.	* (minutes) = every minute of the hour. * (days) = every day.
*/N	Every N th item.	*/15 (minutes) = every 15th minute of the hour. */3 (days) = every 3rd day. */3 (months) = every 3rd month.
Comma and hyphen/dash	Each stated item (comma) Each item in a range (hyphen/dash).	1,31 (minutes) = on the 1st and 31st minute of the hour. 1-3,31 (minutes) = on the 1st to 3rd minutes inclusive, and the 31st minute, of the hour. mon-fri (days) = every Monday to Friday inclusive (every weekday). mar,jun,sep,dec (months) = every March, June, September, December.

You can specify days of the month or days of the week.

TrueNAS lets users create flexible schedules using the available options. The table below has some examples:

Desired schedule	Values to enter
3 times a day (at midnight, 08:00 and 16:00)	months=*; days=*; hours=0/8 or 0,8,16; minutes=0 (Meaning: every day of every month, when hours=0/8/16 and minutes=0)
Every Monday/Wednesday/Friday, at 8.30 pm	months=*; days=mon,wed,fri; hours=20; minutes=30
1st and 15th day of the month, during October to June, at 00:01 am	months=oct-dec,jan-jun; days=1,15; hours=0; minutes=1
Every 15 minutes during the working week, which is 8am - 7pm (08:00 - 19:00) Monday to Friday	Note that this requires two tasks to achieve: (1) months=*; days=mon-fri; hours=8-18; minutes=*/15 (2) months=*; days=mon-fri; hours=19; minutes=0 We need the second scheduled item, to execute at 19:00, otherwise we would stop at 18:45. Another workaround would be to stop at 18:45 or 19:45 rather than 19:00.

Setting **Only Replicate Snapshots Matching Schedule** restricts the replication to only replicate those snapshots created at the same time as the replication schedule.



Unlocking a Replicated Encrypted Dataset or Zvol Without a Passphrase

TrueNAS SCALE users should either replicate the dataset/Zvol without properties to disable encryption at the remote end or construct a special json manifest to unlock each child dataset/zvol with a unique key.

Method 1: Construct JSON Manifest

1. Replicate every encrypted dataset you want to replicate with properties.
2. Export key for every child dataset that has a unique key.
3. For each child dataset construct a proper json with poolname/datasetname of the destination system and key from the source system like this:
{"tank/share01": "57112db4be777d93fa7b76138a68b790d46d6858569bf9d13e32eb9fda72146b"}
4. Save this file with the extension .json.
5. On the remote system, unlock the dataset(s) using properly constructed json files.

Method 2: Replicate Encrypted Dataset/zvol Without Properties

Uncheck properties when replicating so that the destination dataset will not be encrypted on the remote side and will not require a key to unlock.

1. Go to **Data Protection** and click **ADD** in the *Replication Tasks* window.
2. Click *Advanced Replication Creation*.
3. Fill out the form as needed and make sure *Include Dataset Properties* is **NOT** checked.
4. Click **Save**.

Method 3: Replicate Key Encrypted Dataset/zvol

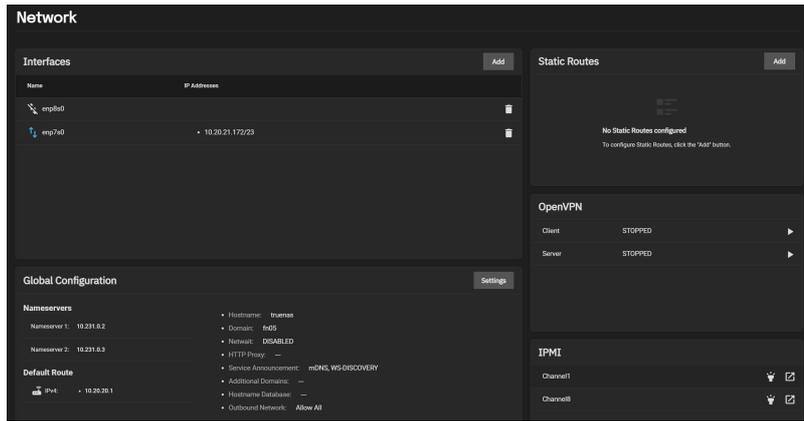
Check **Full Filesystem Replication** so that the destination dataset will use the exported Encryption key from the source pool/dataset to unlock.

1. Go to **Storage -> pool/root dataset**. Click **!** and select **Export Key**.
2. Download the key, open the text file, and copy the Key code.
3. Go to **Data Protection** and click **ADD** in the *Replication Tasks* window.
4. Click *Advanced Replication Creation*.
5. Fill out the form as needed and make sure to enable *Full Filesystem Replication*.
6. Click **Save**.
7. On the receiving pool/dataset:
 - Click **!** next to pool/dataset and select *Unlock*.
 - Unset *Unlock with Key file*.
 - Paste the Key Code into Dataset Key. (if there is a space character at the end of the key - delete the space.)
 - Click **Save**.
 - Click *Continue*.

7 - Network

The SCALE **Network** screen has network configuration and settings options for active interfaces, static routes, and the global configuration.

The **Network** screen also displays OpenVPN information and IPMI channels.



Each networking configurable is a separate widget on the overview screen. Click the buttons or an existing widget entry to view a side panel with configuration options.

Select a networking section in the box below to see more details about specific configuration options.

Networking Tour Video

This video demonstrates configuring networking settings.



Video Player is loading.
Video URL: <https://www.citenas.com/docs/files/scaleangelfishgeneralnetworktour.mp4>

Play Video

Play

Mute

Interfaces

Interfaces

Current Time 0:00

Duration 0:44

The **Interfaces** section displays network port names and IP addresses, as well as their upload/download rates.

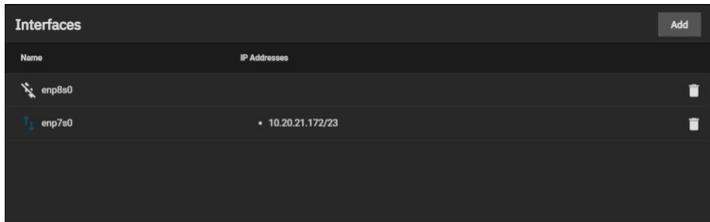
Loaded: 100.00%

Stream Type LIVE

Seek to live, currently behind liveLIVE

Remaining Time -0:44

1x



Click on an interface to edit it, click the **captions off, selected** icon next to the interface to delete it, or click **Add** to add a new one.

Why should I use different interface types?

LAGG (Link Aggregation)

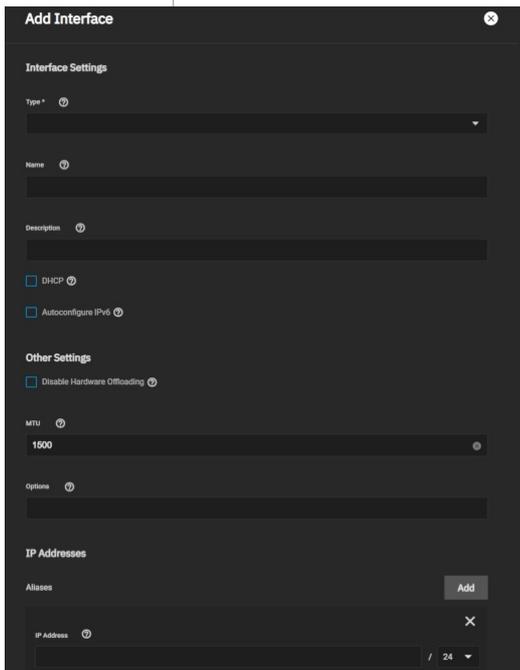
You should use LAGG if you want to optimize multi-user performance, balance network traffic, or have network failover protection.

For example, Failover LAGG prevents a network outage by dynamically reassigning traffic to another interface when one physical link (a cable or NIC) fails.

Network Bridge This is a modal window.

You should use a Bridge if you want to enable communication between two networks and provide a way for them to work as a single network.

For example, bridges can serve IPs to multiple VMs on one interface, which allows your VMs to be on the same network as the host.



Interface Settings

Setting	Description
Type	Choose the type of interface. Bridge creates a logical link between multiple networks. Link Aggregation combines multiple network connections into a single interface. A Virtual LAN (VLAN) partitions and isolates a segment of the connection. Read-only when editing an interface.
Name	Enter a name for the interface. Use the format bondX, vlanX, or brX where X is a number representing a non-parent interface. Read-only when editing an interface.
Description	Enter a description of the interface.
DHCP	Select to enable DHCP. Leave checkbox cleared to create a static IPv4 or IPv6 configuration. Only one interface can be configured for DHCP.
Autoconfigure IPv6	Set to automatically configure the IPv6 address with rtsol(8) . Only one interface can be configured this way.

Other Settings

Setting	Description
Disable Hardware Offloading	Turn off hardware offloading for network traffic processing.
MTU	Maximum Transmission Unit, the largest protocol data unit that can be communicated. The largest workable MTU size varies with network interfaces and equipment. 1500 and 9000 are standard Ethernet MTU sizes. Leaving blank restores the field to the default value of 1500.
Options	Enter additional parameters from ifconfig(8) .

WARNING: Disabling hardware offloading can reduce network performance. We only recommend disabling hardware offloading when the interface is managing jails, plugins, or virtual machines.

IP Addresses

The **IP Address** section lets users define an alias for the interface on the TrueNAS controller. The alias can be an IPv4 or IPv6 address.

Users may also select how many bits are a part of the network address.

Global Configuration

Global Configuration



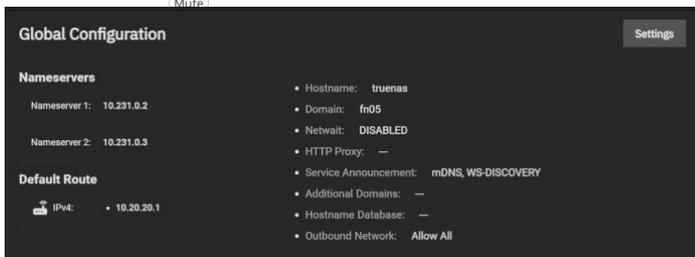
Video Player is loading.
Video URL: <https://www.truenas.com/docs/files/scaleangelfishstaticipglobalnetworking.mp4>

Play Video

Play

Mute

The **Global Configuration** section has all the general TrueNAS networking settings *not* specific to any interfaces.



Chapters

Can I configure these options elsewhere?

Users can configure many of these interface, DNS, and gateway options in the [Console Setup Menu](#). Be sure to check both locations when troubleshooting network connectivity issues.

Descriptions

Disruptive Change

- descriptions off, selected

You can lose your TrueNAS connection if you change the network interface that the web interface uses! You might need command line knowledge or physical access to the TrueNAS system to fix misconfigured network settings.

- captions off, selected

Audio Track

- Unknown, selected

Fullscreen

This is a modal window.

Beginning of dialog window. Escape will cancel and close the window.

Text

Color Transparency

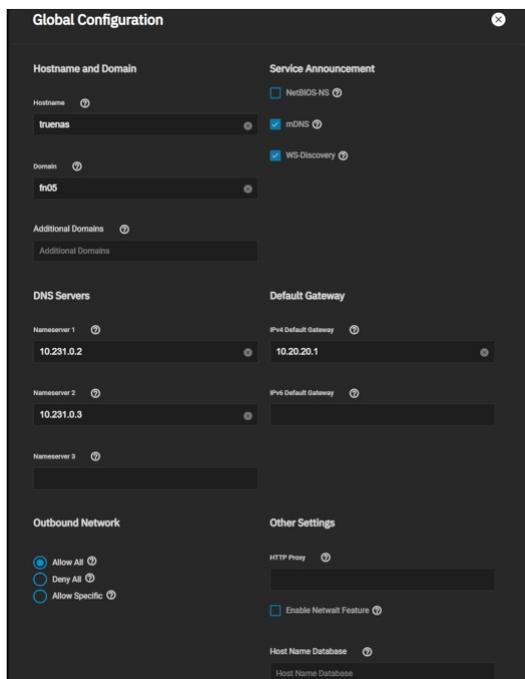
Background

Color Transparency

Window

Color Transparency

Font Size



Hostname and Domain

Many of these fields have default values, but users can change them to meet local network requirements.

TrueNAS displays the **Hostname** and **Domain** in the **Dashboard System Information** window.

Note: The **Global Configuration** window only displays some options when the appropriate hardware is present.

Setting	Description
Hostname	System hostname.
Hostname (TrueNAS Controller 2)	Host name of second TrueNAS controller (for HA only). Upper and lower case alphanumeric, ., and - characters are allowed.
Hostname (Virtual)	Virtual host name. When using a virtual host, this is also used as the Kerberos principal name. Enter the fully qualified hostname plus the domain name. Upper and lower case alphanumeric, ., and - characters are allowed.
Domain	System domain name, like example.com
Additional Domains	Additional domains to search. Separate entries by pressing Enter . Adding search domains can cause slow DNS lookups

Service Announcement

Setting	Description
NetBIOS-NS	Legacy NetBIOS name server. Advertises the SMB service NetBIOS Name. Can be required for legacy SMB1 clients to discover the server. When advertised, the server appears in Network Neighborhood .
mDNS	Multicast DNS. Uses the system hostname to advertise enabled and running services. For example, this controls if the server appears under Network on MacOS clients.
WS-Discovery	Uses the SMB Service NetBIOS name to advertise the server to WS-Discovery clients. This causes the computer to appear in the Network Neighborhood of modern Windows OSes.

DNS Servers

Setting	Description
Nameserver 1	Primary DNS server.
Nameserver 2	Secondary DNS server.
Nameserver 3	Third DNS server.

Default Gateway

Setting	Description
IPv4 Default Gateway	Enter an IPv4 address. This overrides the default gateway provided by DHCP.
IPv6 Default Gateway	Enter an IPv6 address. This overrides the default gateway provided by DHCP.

Outbound Network

Setting	Description
Allow All	Any system service can communicate externally.
Deny All	This system cannot communicate externally.
Allow Specific	Define the system services that are allowed to communicate externally. All other external traffic is restricted.

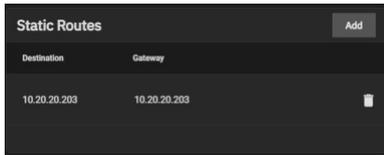
Other Settings

Setting	Description
HTTP Proxy	When using a proxy, enter the proxy information for the network in the format <code>http://my.proxy.server:3128</code> or <code>http://user:password@my.proxy.server:3128</code> .
Enable Netwait Feature	Delays the start of network services until pings return from the IP addresses added to the Netwait IP List .
Netwait IP List	Only appears when Enable Netwait Feature checkbox is selected. Enter a list of IP addresses to ping . Separate entries by pressing Enter . Each address is tried until one is successful or the list is exhausted. Leave empty to use the default gateway.
Host Name Database	Additional hosts to append to <code>/etc/hosts</code> . Separate entries by pressing Enter . Use the format <code>IP_address space hostname</code> where multiple hostnames can be used if separated by a space. Hosts defined here are still accessible by name even when DNS is not available. See hosts for additional information.

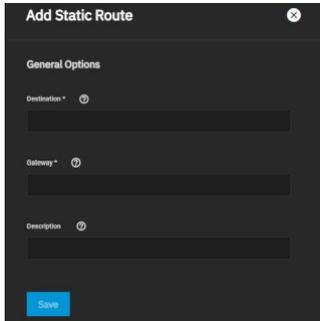
Static Routes

Static Routes

TrueNAS administrators can use the **Static Routes** section to manually enter routes so the router can send packets to a destination network.



TrueNAS does not have defined static routes by default. If you need a static route to reach portions of the network, add the route by going to **Network** and clicking **Add** in the **Static Routes** window.



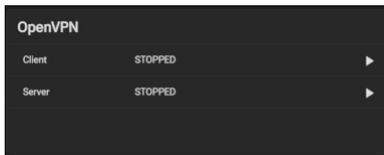
Setting	Description
Destination	Use the format <i>A.B.C.D/E</i> where <i>E</i> is the CIDR mask.
Gateway	Enter the IP address of the gateway.
Description	Notes or identifiers describing the route.

OpenVPN

OpenVPN

A virtual private network (VPN) is an extension of a private network over public resources. VPNs allow clients to securely connect to a private network even when remotely using a public network.

TrueNAS provides **OpenVPN** as a system-level service for VPN server or client functionality. TrueNAS can act as a primary VPN server that gives remote clients access to data stored on the system using a single TCP or UDP port. Alternately, TrueNAS can integrate into a private network, even when the system is in a separate physical location or only has access to publicly visible networks.



Before configuring TrueNAS as either an OpenVPN server or client, you need an existing public key infrastructure (PKI) with [Certificates](#) and [Certificate Authorities](#) created in or imported to TrueNAS.

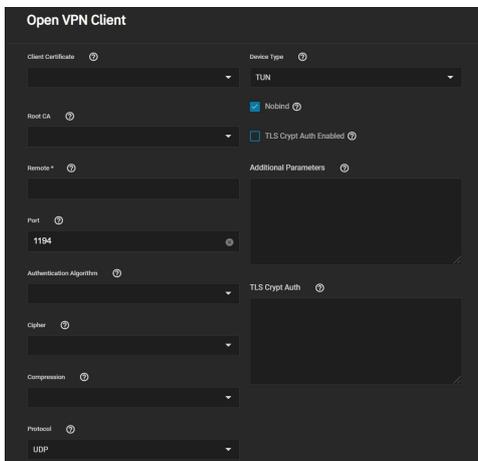
What does a PKI do?

A PKI allows TrueNAS to authenticate with clients or servers by confirming a valid master Certificate Authority (CA) signed the network credentials. See the [OpenVPN PKI Overview](#) to read more about the OpenVPN required PKI.

To configure OpenVPN (server or client) on TrueNAS, select the networking credentials, set the connection details, and choose additional security or protocol options.

OpenVPN Client

Go to **Network** and click **Client** in the **OpenVPN** window to configure the OpenVPN client.

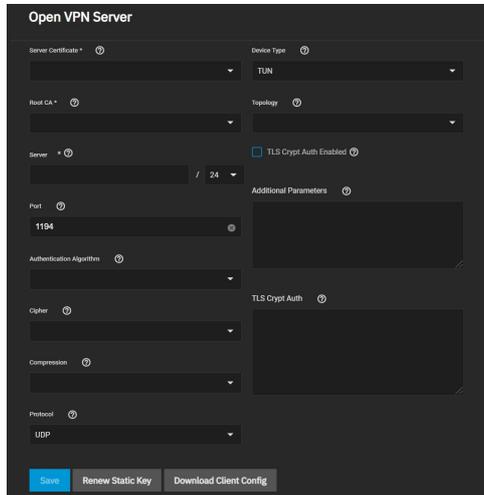


Setting	Description
Client Certificate	Choose a valid client certificate that exists on this system and is not revoked. Find more about generating certificates and CAs for OpenVPN here .
Root CA	Choose the root Certificate Authority that you used to sign the client and server certificates. Find more about generating certificates and CAs for OpenVPN here .
Remote	A valid IP address or domain name to which OpenVPN connects.
Port	Enter a port number to use for the connection.
Authentication Algorithm	Choose an algorithm to authenticate packets.

Cipher	Choose a cipher algorithm to encrypt data channel packets.
Compression	Choose a compression algorithm.
Protocol	Choose the protocol to use when connecting with the remote system.
Device Type	Choose a virtual network interface. More information is found here .
Nobind	Enable to prevent binding to local address and port. Must be enabled if OpenVPN client and server are to run concurrently.
TLS Crypt Auth Enabled	Enable/disable TLS Web Client Authentication.
Additional Parameters	Additional parameters.
TLS Crypt Auth	Provide static key for authentication/encryption of all control channel packets when <code>tls_crypt_auth_enabled</code> is enabled.

OpenVPN Server

Go to **Network** and click **Server** in the **OpenVPN** window to configure the OpenVPN server.



Setting	Description
Server Certificate	Choose a valid client certificate which exists on this system and is not revoked. Find more about generating certificates and CAs for OpenVPN here .
Root CA	Choose the root Certificate Authority that you used to sign the client and server certificates. Find more about generating certificates and CAs for OpenVPN here .
Server	Enter the IP address and netmask of the server.
Port	Enter a port number to use for the connection.
Authentication Algorithm	Choose an algorithm to authenticate packets.
Cipher	Choose a cipher algorithm to encrypt data channel packets.
Compression	Choose a compression algorithm.
Protocol	Choose the protocol to use when connecting with the remote system.
Device Type	Choose a virtual network interface. More information is found here .
Topology	Configure virtual addressing topology when running in TUN mode. (TAP mode always uses a SUBNET topology.)
TLS Crypt Auth Enabled	Enable/disable TLS Web Client Authentication.
Additional Parameters	Additional parameters.
TLS Crypt Auth	Provide static key for authentication/encryption of all control channel packets when <code>tls_crypt_auth_enabled</code> is enabled.

TUN Device Type

If you choose the **TUN** as the **Device Type**, you can select a virtual addressing **Topology** for the server:

- **NET30**: Use one /30 subnet per client in a point-to-point topology. Designed for use when connecting clients are Windows systems.
- **P2P**: Point-to-point topology that points the local server and remote client endpoints to each other. P2P gives each client one IP address. We only recommend P2P when none of the clients are Windows systems.
- **SUBNET**: the interface uses an IP address and subnet. SUBNET gives each client one IP address. Windows clients require the **TAP-Win32 driver** version 8.2 or newer. **TAP** devices always use the **SUBNET** for **Topology**.

TrueNAS automatically applies the **Topology** selection to any connected clients.

TLS Crypt Auth

When users enable **TLS Crypt Auth Enabled**, TrueNAS generates a static key for the **TLS Crypt Auth** field after saving the options. To change this key, click **Renew Static Key**. Clients connecting to the server require the static key. TrueNAS stores the keys in the system database and automatically includes them in a generated client config file. We always recommend users back up keys in a secure location.

After configuring and saving your OpenVPN server, generate client configuration files to importing to OpenVPN client systems connecting to the server. You need the certificate from the client system already imported onto the system. To generate the configuration file, click **Download Client Config** and select the **Client Certificate**.

Common Options (Client or Server)

Many of the fields for configuring an OpenVPN server or client are identical.

The **Additional Parameters** field manually sets any of the core OpenVPN config file options. See the OpenVPN [Reference Manual](#) for option descriptions.

Connection Settings

- **Root CA**: The Certificate Authority (CA) must be the root CA TrueNAS used to sign the client and server certificates.
- **Port**: This is the port that the OpenVPN connection uses.
- **Compression**: Choose a compression algorithm for traffic. Leave the field empty for data to be sent uncompressed. **LZO** is a standard compression algorithm that is backward-compatible with previous (pre-2.4) versions of OpenVPN. **LZ4** is a newer option that is typically faster with fewer system resources required.
- **Protocol**: Choose between **UDP** or **TCP** protocols for OpenVPN. UDP sends packets in a continuous stream while TCP sends packets sequentially. UDP is generally faster and less strict about dropped packets than TCP. To force the connection to be IPv4 or IPv6, choose one of the 4 or 6 **UDP** or **TCP** options.
- **Device Type**: use a **TUN** or **TAP** virtual networking device and layer with OpenVPN. The device type must be identical between the OpenVPN server and any clients.

Security Options

Because using a VPN involves connecting to a private network while still sending data over less secure public resources, OpenVPN includes several security options. While not required, these security options help protect the data TrueNAS sends into or out of the private network.

- **Authentication Algorithm:** Validates packets that TrueNAS sends over the network connection. Your network environment might require a specific algorithm. If not, SHA1 HMAC is a good standard algorithm to use.
- **Cipher:** Encrypts data packets sent through the connection. While not required, choosing a cipher can increase connection security. You might need to verify which ciphers your networking environment requires. If there are no specific cipher requirements, AES-256-GCM is a good default choice.
- **TLS Encryption:** When **TLS Crypt Auth Enabled** is selected, TrueNAS encrypts all TLS handshake messages to add another layer of security. TLS Encryption requires a static key that the OpenVPN server and clients share.

Service Activation

When finished configuring the server or client service, click **Save**. Start the service by clicking the play button next to it in the **OpenVPN** window.

You may also start the service by going to **System Settings > Services** and clicking the **State** toggle. Setting **Start Automatically** starts the service when TrueNAS completes booting and runs the network and data pools.

IPMI

IPMI

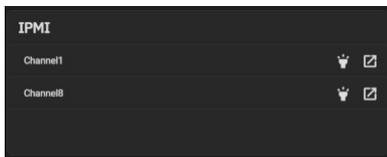
IPMI (Intelligent Platform Management Interface) requires compatible hardware! Refer to your hardware documentation to determine if the TrueNAS web interface has IPMI options.

Many [TrueNAS Storage Arrays](#) have a built-in out-of-band management port that provides side-band management should the system become unavailable through the web interface.

IPMI allows users to check the log, access the BIOS setup, and boot the system without physical access. IPMI also enables users to remotely access the system to assist with configuration or troubleshooting issues.

Some IPMI implementations require updates to work with newer versions of Java. See [here](#) for more information.

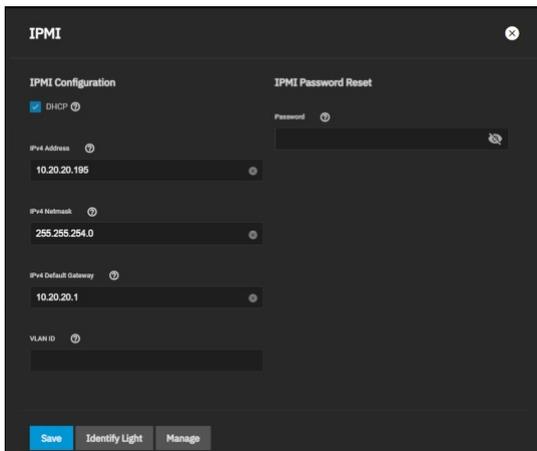
IPMI is configured in **Network > IPMI**. The IPMI configuration screen provides a shortcut to the most basic IPMI configuration.



The **IPMI** window displays the available IPMI channels. The **Identify Light** button (flashlight) lets users select a duration for the system's IPMI to flash so they can identify it. The **Manage** button (square with an outward-pointing arrow) opens the IPMI manager in a new browser tab.

IPMI Configuration

Click the channel you wish to edit to open the configuration form.



Setting	Description
DHCP	Use DHCP. Clear checkbox to manually configure a static IPv4 connection.
IPv4 Address	Static IPv4 address of the IPMI web interface.
IPv4 Netmask	Subnet mask of the IPv4 address.
IPv4 Default Gateway	Enter the default gateway of the IPv4 connection.
VLAN ID	Enter the VLAN identifier if the IPMI out-of-band management interface is not on the same VLAN as management networking.
Password	Enter the password used to connect to the IPMI interface from a web browser. The maximum length accepted in the UI is 20 characters, but different hardware might require shorter passwords.
Identify Light	Lets users select a duration for the system's IPMI light to flash on the compatible connected hardware.
Manage	Opens the IPMI manager in a new browser tab.

IPMI Options

After saving the configuration, users can access the IPMI interface using a web browser and the IP address specified in **Network > IPMI**. The management interface prompts for login credentials. Refer to your IPMI device documentation to learn the default administrator account credentials.

After logging in to the management interface, users can change the default administrative user name and create additional IPMI users. IPMI utility appearance and available functions vary by hardware.

8 - Credentials

SCALE Credential options are collected in this section of the UI and organized into a few different screens:

- **Local Users** allows those with permissions to add, configure, and delete users on the system. There are options to search for keywords in usernames, display or hide user characteristics, and toggle whether the system shows built-in users.
- **Local Groups** allows those with permissions to add, configure, and delete user groups on the system. There are options to search for keywords in group names, display or hide group characteristics, and toggle whether the system shows built-in groups.
- **Directory Services** contains options to edit directory domain and account settings, set up Idmapping, and configure access and authentication protocols. Specific options include configuring Kerberos realms and key tables (keytab), as well as setting up LDAP validation.
- **Backup Credentials** stores credentials for cloud backup services, SSH Connections, and SSH Keypairs. Users can set up backup credentials with cloud and SSH clients to back up data in case of drive failure.
- **Certificates** contains all the information for certificates, certificate signing requests, certificate authorities, and DNS-authenticators. TrueNAS comes equipped with an internal, self-signed certificate that enables encrypted access to the web interface, but users can make custom certificates for authentication and validation while sharing data.
- **2FA** allows users to set up Two-Factor Authentication for their system. Users can set up 2FA, then link the system to an authenticator app (such as Google Authenticator, LastPass Authenticator, etc.) on a mobile device.

Ready to get started? Choose a topic or article from the left-side Navigation pane. Click the < symbol to expand the menu to show the topics under this section.

8.1 - Local Users

- [Creating User Accounts](#)

In TrueNAS, user accounts allow flexibility for accessing shared data. Typically, administrators create users and assign them to [groups](#). Doing so makes tuning permissions for large numbers of users more efficient.

Only the *root* user account can log in to the TrueNAS web interface.

When the network uses a directory service, import the existing account information using the instructions in [Directory Services](#).

Using [Active Directory](#) requires setting Windows user passwords in Windows.

To see user accounts, go to **Credentials > Local Users**.

Username	UID	Builtin	Full Name
root	0	true	root
user1	1000	false	user1
user2	1001	false	user2

TrueNAS hides all built-in users (except root) by default. Click the **SHOW** button, then click **SHOW** to see all built-in users.

Creating User Accounts

Tutorial Video

This short video demonstrates adding a local user.



To create a new user, click **Add**.

Play Video

Play

Mute

Current Time 0:00 / Duration 0:47

Loaded: 100.00%

Stream Type LIVE

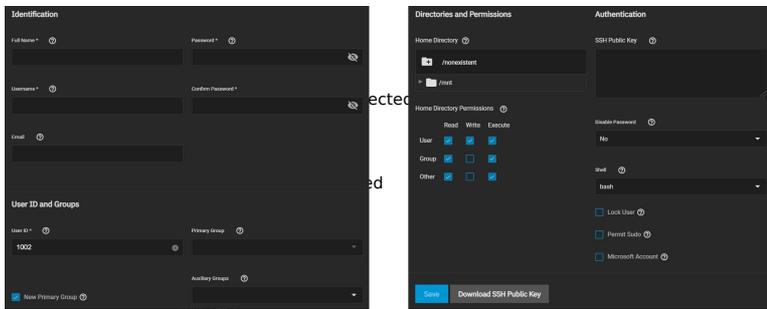
Seek to live, currently behind liveLIVE

Remaining Time -0:47

1x

Playback Rate

Chapters



TrueNAS lets users configure four different user account traits. This is a modal window.

Identification Beginning of dialog window. Escape will cancel and close the window.

Identification Text
Enter the user full name in the **Full Name** field. TrueNAS suggests a simplified name in **Username** derived from the **Full Name**, but you can override it with your own choice.

Background
You can also assign a user account email address in the **Email** field.

Set and confirm a password
Window

User ID and Groups Color: Black Transparency: Transparent

User ID and Groups Size

100%
Next, you must set a user ID. TrueNAS suggests a user ID starting at **1000**, but you can change it if you wish. We recommend using an ID of 1000 or greater for non-built-in users.

Text Edge Style

By default, TrueNAS creates a new primary group with the same name as the user. To add the user to an existing primary group instead, clear the **New Primary Group** checkbox and select a group from the **Primary Group** drop-down list. You can add the user to more groups using the **Auxiliary Groups** drop-down list.

Directories and Permissions Proportional Sans-Serif

Directories and Permissions Does not create sub-directories to the default values Done

When creating a user, the home directory path is set to /nonexistent, which does not create a home directory for the user. To set a user home directory, select a path using the file browser. If the directory exists and matches the user name, TrueNAS sets it as the user home directory. When the path does not end with a sub-directory matching the user name, TrueNAS creates a new sub-directory. TrueNAS shows the path to the user home directory when editing a user.

Close Modal Dialog
end of dialog window

You can set the home directory permissions directly under the file browser. You cannot change TrueNAS default user account permissions.

Authentication

Authentication

You can assign a public SSH key to a user for key-based authentication by pasting the *public* key into the **SSH Public Key** field. If you are using an SSH public key, always keep a backup of the key. Click **Download SSH Public Key** to download the pasted key as a .txt file.

By default, **Disable Password** is **No**.

Setting **Disable Password** to **Yes** disables several options:

- The **Password** field becomes unavailable, and TrueNAS removes any existing password from the account.
- The **Lock User** and **Permit Sudo** options disappear.
- The account is restricted from password-based logins for services like SMB shares and SSH sessions.

You can set a specific [shell](#) for the user from the **Shell** drop-down:

Shell	Description
csh	C shell for UNIX system interactions.
sh	Bourne shell
tcsh	Enhanced C shell that includes editing and name completion.
bash	Bourne Again shell for the GNU operating system.
ksh93	Korn shell that incorporates features from both <i>csh</i> and <i>sh</i> .
mksh	MirBSD Korn Shell
rbash	Restricted bash
rzsh	Restricted zsh
scponly	scponly restricts the user's SSH usage to only the <i>scp</i> and <i>sftp</i> commands.
zsh	Z shell
git-shell	restricted git shell
nologin	Use when creating a system account or to create a user account that can authenticate with shares but that cannot log in to the TrueNAS system using <i>ssh</i> .

Setting **Lock User** disables all password-based functionality for the account until you unset it.

Permit Sudo allows the account to act as the system administrator using the *sudo* command. Leave it disabled for better security.

If the user accesses TrueNAS data using *Windows 8* or newer, set **Microsoft Account** to enable those systems additional authentication methods.

By default, **Samba Authentication** is enabled. This allows using the account credentials to access data shared with [SMB](#).

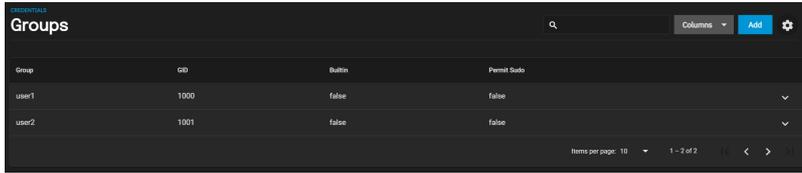
8.2 - Local Groups

- [View Existing Groups](#)
- [Add a New Group](#)
- [Group Member Management](#)

TrueNAS offers groups as an efficient way to manage permissions for many similar user accounts. See [Users](#) for managing users. The interface lets you manage UNIX-style groups. If the network uses a directory service, import the existing account information using the instructions in [Active Directory](#).

View Existing Groups

To see saved groups, go to **Credentials > Local Groups**.

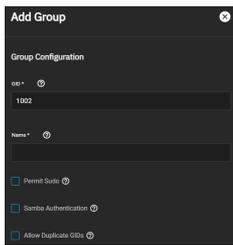


By default, TrueNAS hides the system's built-in groups. To see built-in groups, click

then click *Show*.

Add a New Group

To create a group, go to **Credentials > Local Groups** and click *Add*.



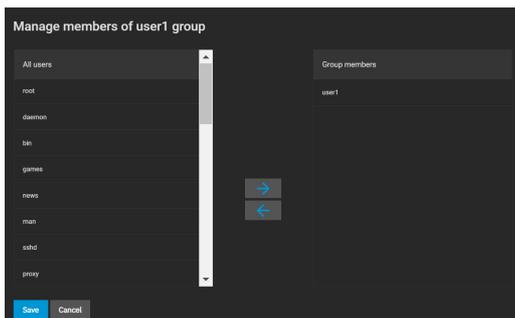
Setting	Description
GID	The Group ID (<i>GID</i>) is a unique number TrueNAS uses to identify a Unix group. Enter a number above 1000 for a group with user accounts (you cannot change the <i>GID</i> later). If a system service uses a group, the group ID must match the service's default port number.
Name	Group name cannot begin with a hyphen (-). Group name cannot contain a space, tab, or any of these characters: , : + & # % ^ () ! @ ~ * ? < > =. You may only use \$ as a username's last character.
Permit Sudo	Allow group members to use sudo . TrueNAS prompts group members for their password when using sudo. Leave <i>Permit Sudo</i> disabled for better security.
Samba Authentication	Set to let Samba permissions and authentication use group.
Allow Duplicate GIDs	Allows more than one group to have the same group ID. We do not recommend enabling this.

Group Member Management

Register user accounts to a group to simplify permissions and access to large numbers of user accounts. To manage group membership, go to **Credentials > Local Groups**, click the

next to a group, and click

Members:



To add user accounts to the group, select them in All users and click →. Select multiple users by holding CTRL while clicking each entry.

8.3 - Directory Services

The SCALE Directory Services section contains options to edit directory domain and account settings, set up Idmapping, and configure authentication and authorization services in TrueNAS SCALE.

Choosing Active Directory or LDAP

When setting up directory services in TrueNAS, you can connect TrueNAS to either an Active Directory or an LDAP server.

Active Directory

The Active Directory (AD) service shares resources in a Windows network. AD provides authentication and authorization services for the users in a network, eliminating the need to recreate the user accounts on TrueNAS.

Once joined to an AD domain, you can use domain users and groups in local ACLs on files and directories. You can also set up shares to act as a file server.

Joining an AD domain also configures the Privileged Access Manager (PAM) to let domain users log on via SSH or authenticate to local services.

Users can configure AD services on Windows or Unix-like operating systems running [Samba version 4](#).

To configure a connection, you will need to know the Active Directory domain controller's domain and that system's account credentials.

Preparation

Users can take a few steps before configuring Active Directory to ensure the connection process goes smoothly.

Verify Name Resolution

To confirm that name resolution is functioning, go to **System Settings > Shell** and use `ping` to check the connection to the AD domain controller.

```
truenas# ping ad02.lab.ixsystems.com
ping: ad02.lab.ixsystems.com (10.215.5.200): 56 data bytes
64 bytes from 10.215.5.200: icmp_seq=0 ttl=126 time=0.800 ms
64 bytes from 10.215.5.200: icmp_seq=1 ttl=126 time=0.933 ms
^C
--- ad02.lab.ixsystems.com ping statistics ---
2 packets transmitted, 2 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 0.800/0.933 ms
truenas#
```

When TrueNAS sends and receives packets without loss, the connection is verified. Press `ctrl + c` to cancel the ping.

Another option is to use `host -t srv _ldap._tcp.domainname.com` to check the network's SRV records and verify DNS resolution.

The ping failed!

If the ping fails, go to **Network** and click *Settings* in the *Global Configuration* window. Update the *DNS Servers* and *Default Gateway* settings so the connection to your Active Directory Domain Controller can start. Use more than one *Nameserver* for the AD domain controllers so DNS queries for requisite SRV records can succeed. Using more than one *Nameserver* helps maintain the AD connection whenever a domain controller becomes unavailable.

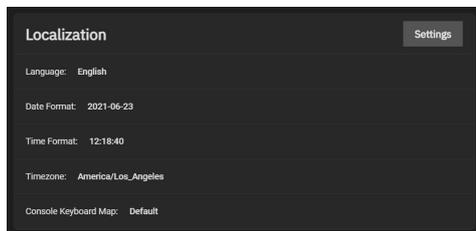
Time Synchronization

Active Directory relies on the time-sensitive [Kerberos](#) protocol. TrueNAS adds the AD domain controller with the [PDC Emulator FSMO Role](#) as the preferred NTP server during the domain join process. If your environment requires something different, go to **System Settings > General** and add or edit a server in the *NTP Servers* window.

The local system time cannot be out of sync by more than **five (5) minutes** with the AD domain controller time in a default AD environment. Use an external time source when configuring a virtualized domain controller. TrueNAS creates an alert if the system time gets out of sync with the AD domain controller time.

TrueNAS has a few options to ensure both systems are synchronized:

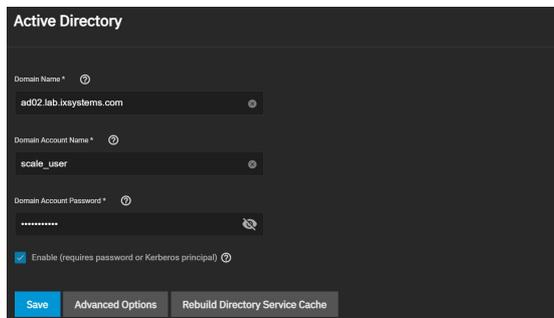
1. Go to **System Settings > General** and click *Settings* in the *Localization* window to ensure the *Timezone* matches the AD Domain Controller.



2. Set either localtime or universal time in the system BIOS.

Connect to the Active Directory Domain

To connect to Active Directory, click *Settings* in the *Active Directory* window and enter the *AD Domain Name* and account credentials. Set *Enable* to attempt to join the AD domain immediately after saving the configuration.



TrueNAS offers advanced options for fine-tuning the AD configuration, but the preconfigured defaults are generally suitable.

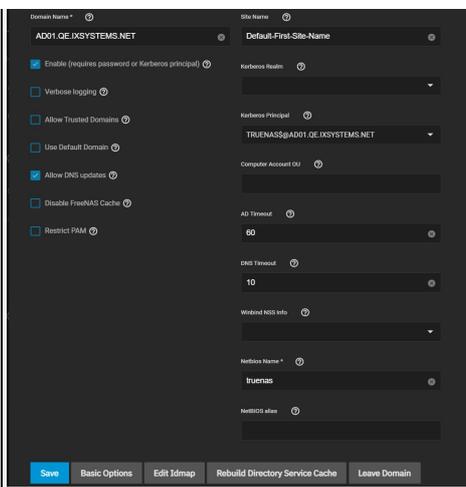
I don't see any AD information!

TrueNAS can take a few minutes to populate the Active Directory information after configuration. To check the AD join progress, open the **Task Manager** in the upper-right corner. TrueNAS displays any errors during the join process in the **Task Manager**.

When the import is complete, AD users and groups become available while configuring basic dataset permissions or an ACL with TrueNAS cache enabled (enabled by default).

Joining AD also adds default Kerberos realms and generates a default `AD_MACHINE_ACCOUNT` keytab. TrueNAS automatically begins using this default keytab and removes any administrator credentials stored in the TrueNAS configuration file.

Advanced Options



Setting	Description
Verbose logging	Logs attempts to join the domain in <code>/var/log/messages</code> .
Allow Trusted Domains	When set, usernames do not include a domain name. Unset to prepend domain names to user names. Unsetting this option prevents username collisions when there are identical usernames across multiple domains.
Use Default Domain	Unset to prepend the domain name to the username and prevent name collisions when using Allow Trusted Domains with the same username across multiple domains.
Allow DNS Updates	Enables Samba to do DNS updates when joining a domain.
Disable FreeNAS Cache	Disables caching AD users and groups, which can help when unable to bind to a domain with a lot of users or groups.
Restrict PAM	Restricts SSH access to <code>BUILTIN\Administrators</code> members in certain circumstances.
Site Name	Enter the relative distinguished name of the site object in the Active Directory.
Kerberos Realm	Select an existing realm from <i>Kerberos Realms</i> .
Kerberos Principal	Select the location of the principal in the keytab created in Directory Services > Kerberos Keytabs.
Computer Account OU	The OU that creates new computer accounts. TrueNAS reads the OU string from top to bottom without RDNs. Uses forward slashes (/) as delimiters, like <code>Computers/Servers/NAS</code> . Use backslashes (\) to escape characters but not as a separator. TrueNAS interprets backslashes at multiple levels, so you might have to use several for them to work. When this field is blank, TrueNAS creates new computer accounts in the AD default OU.
AD Timeout	Number of seconds before timeout. To view the AD connection status, open the interface Task Manager.
DNS Timeout	Number of seconds before a timeout. Increase this value if AD DNS queries time out.
Winbind NSS Info	Choose the schema to use when querying AD for user/group info. <code>rfc2307</code> uses the Windows 2003 R2 schema support, <code>sfu</code> is for Service For Unix 3.0 or 3.5, and <code>sfu20</code> is for Service For Unix 2.0.
Netbios Name	Netbios Name of this NAS. This name must differ from the Workgroup name and be no greater than 15 characters.
NetBIOS alias	Alternative names (no greater than 15 characters) that SMB clients can use when connecting to this NAS. Can be no greater than 15 characters.
Edit Idmap	Navigates to Directory Services > Idmap so the user can edit the Active Directory's Idmap
Leave Domain	Disconnects the TrueNAS system from the Active Directory.

Troubleshooting

Resync the Cache

If the cache becomes out of sync or fewer users than expected are available in the permissions editors, resync it by clicking *Settings* in the *Active Directory* window and selecting *Rebuild Directory Service Cache*.

If you are using Windows Server with 2008 R2 or older, try creating a **Computer** entry on the Windows server Organizational Unit (OU).

When creating the entry, enter the TrueNAS hostname in the name field and make sure it matches the:

- Hostname:** Go to **Network** and find *Hostname* in the *Global Configuration* window.
- NetBIOS alias:** Go to **Credentials > Directory Services** and click *Settings* in the *Active Directory* window. Click *Advanced Options* and find the *NetBIOS alias*.

Shell Commands

You can go to **System Settings > Shell** and enter various commands to get more details about the AD connection and users:

AD current state: `midclt call activedirectory.get_state`.

Connected LDAP server details: `midclt call activedirectory.domain_info | jq`. For example:

```
truenas# midclt call activedirectory.domain_info | jq
{
  "LDAP server": "192.168.1.125",
  "LDAP server name": "DC01.HOMEDOM.FUN",
  "Realm": "HOMEDOM.FUN",
  "Bind Path": "dc=HOMEDOM,dc=FUN",
  "LDAP port": 389,
  "Server time": 1593026080,
  "KDC server": "192.168.1.125",
  "Server time offset": 5,
  "Last machine account password change": 1592423446
}
```

View AD users: `wbinfo -u`.

- Enter `getent passwd DOMAIN\<user>` to see more user details (<user> = desired user name).
- If `wbinfo -u` shows more users than are available when configuring permissions and the TrueNAS cache is enabled, go to **Directory Services**, click *Settings* in the *Active Directory* window, and increase the *AD Timeout* value.

View AD groups: `wbinfo -g`. Enter `getent group DOMAIN\domain\ users` to see more details.

View domains: `wbinfo -m`.

Test AD connection: `wbinfo -t`.

- A successful test shows a message like checking the trust secret for domain YOURDOMAIN via RPC calls succeeded.

Test user connection to SMB share: `smbclient '//0.0.0.0/smbshare -U AD.DOMAIN.COM\user`

- 0.0.0.0 is the server address
- smbshare is the SMB share name
- AD.DOMAIN.COM is the trusted domain
- user is the user account name to authenticate.

Clean Up Active Directory

TrueNAS SCALE requires users to cleanly leave an Active Directory using the *Leave Domain* button under *Advanced Settings* to remove the AD object.

If the AD server moves or shuts down without you using *Leave Domain*, TrueNAS won't remove the AD object, and you will have to clean up the Active Directory.

Go to **Credentials > Directory Services** and click *Show* next to *Advanced Settings*

1. Clean out Kerberos settings by clicking *Settings* in the *Kerberos Settings* window and clearing the *Appdefaults Auxiliary Parameters* and *Libdefaults Auxiliary Parameters* boxes. You may also need to clear out leftover Kerberos Realms and Keytabs by clicking the *Clear* next to remaining entries.
2. Click the *Idmap Active Directory - Primary Domain* entry and clear out the Active Directory settings, then click *CONTINUE* to clear the Idmap cache.
3. Go to **Network** and click *Settings* in the *Global Configuration* window. Remove the Active Directory Nameserver and enter a new one.
4. Ensure all other network settings are correct.
5. Go to **System Settings > Services** and change the workgroup to "WORKGROUP".
6. Go to **Credentials > Directory Services** and edit the Active Directory config to the new domain.
7. Make sure the Kerberos settings and Idmap are correct and that SMB is running.

LDAP

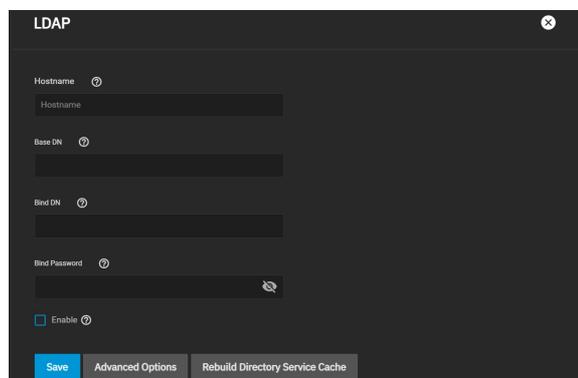
TrueNAS has an [Open LDAP](#) client for accessing the information on an LDAP server. An LDAP server provides directory services for finding network resources such as users and their associated permissions.

Does LDAP work with SMB?

LDAP authentication for SMB shares is disabled unless you have configured and populated the LDAP directory with Samba attributes. The most popular script for performing this task is `smbldap-tools`. The LDAP server must support SSL/TLS, and you must import the certificate for the LDAP server CA. TrueNAS does not support non-CA certificates.

Configuration

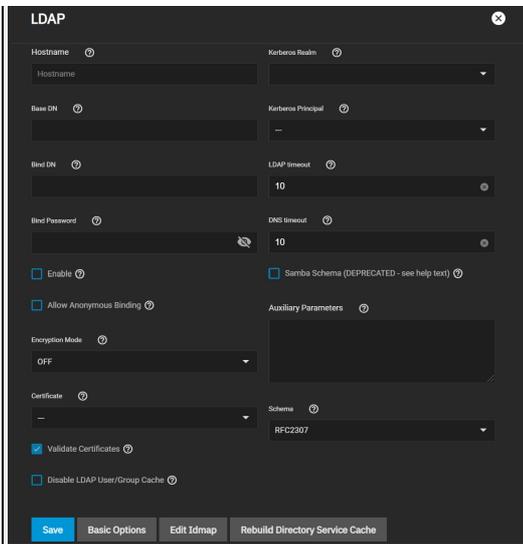
To integrate an LDAP server with TrueNAS, go to **Credentials > Directory Services** and click *Settings* in the *LDAP* window.



Field	Description
Hostname	LDAP server hostnames/IP addresses. Separate entries with Space. You can enter multiple hostnames/IP addresses to create an LDAP failover priority list. If a host does not respond, TrueNAS will try the next host until it establishes a connection.
Base DN	Top level of the LDAP directory tree to be used when searching for resources. Example: <code>dc=test,dc=org</code> .
Bind DN	Administrative account name on the LDAP server. Example: <code>cn=Manager,dc=test,dc=org</code> .
Bind Password	Password for the Bind DN.
Enable	Activates the configuration. Unset to disable the configuration without deleting it. You can re-enable it later without reconfiguring it.

Advanced Configuration

To further modify the LDAP configuration, click *Advanced Options*.



Field	Description
Allow Anonymous Binding	Set for the LDAP server to disable authentication and allow read and write access to any client.
Encryption Mode	Options for encrypting the LDAP connection:

- **OFF**: do not encrypt the LDAP connection.
- **ON**: encrypt the LDAP connection with SSL on port 636.
- **START_TLS**: encrypt the LDAP connection with STARTTLS on the default LDAP port 389. | | Certificate | Certificate to use when performing LDAP certificate-based authentication. To configure LDAP certificate-based authentication, create a Certificate Signing Request for the LDAP provider to sign. TrueNAS does not need a certificate when using username/password or Kerberos authentication. To configure LDAP certificate-based authentication, [create a Certificate Signing Request](#) for the LDAP provider to sign. | | Validate Certificates | Verify certificate authenticity. | | Disable LDAP User/Group Cache | Disable caching LDAP users and groups in large LDAP environments. When caching is disabled, LDAP users and groups do not appear in drop-down menus but are still accepted when manually entered. | | Kerberos Realm | Select an existing realm from *Kerberos Realms*. | | Kerberos Principal | Select the location of the principal in the keytab created in *Kerberos Keytab*. | | LDAP timeout | LDAP timeout in seconds. Increase this value if a Kerberos ticket timeout occurs. | | DNS timeout | DNS timeout in seconds. Increase this value if DNS queries timeout. | | Samba Schema | Only set if you configured the LDAP server with Samba attributes and it requires LDAP authentication for SMB shares. | | Auxiliary Parameters | You can specify additional options for [nscd.conf](#). | | Schema | Schema to use with Samba Schema. | | Edit Idmap | The *Edit Idmap* button takes users to the *Idmap* configuration screen. |

DEPRECATED: Samba Schema support is deprecated in Samba 4.13. We will remove this feature after Samba 4.14. Users should begin upgrading legacy Samba domains to Samba AD domains.

Troubleshooting

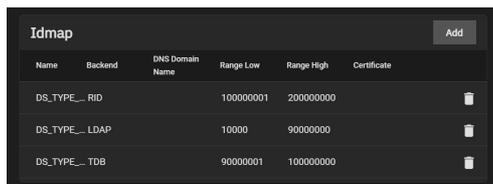
If the cache becomes out of sync or fewer users than expected are available in the permissions editors, resync the cache using the *Rebuild Directory Service Cache*.

Advanced Settings

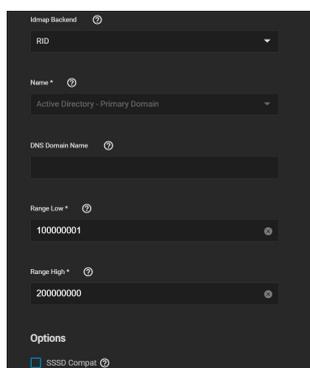
To view *Idmap* and *Kerberos Services*, click *Show* next to *Advanced Settings*.

Idmap

The *Idmap* directory service lets users configure and select a backend to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs. Users must enable the *Active Directory* service to configure and use Identity Mapping (*Idmap*).



Users can click *Add* in the *Idmap* window to configure backends or click on an already existing *Idmap* to edit it.



Field	Description
Idmap Backend	Provides a plugin interface for Winbind to use varying backends to store SID/uid/gid mapping tables. The correct setting depends on the environment you deployed the NAS in.
Name	Enter the pre-Windows 2000 domain name.
DNS Domain Name	DNS name of the domain.
Range Low	Range Low and Range High set the range of UID/GID numbers the IDMap backend translates. If an external credential like a Windows SID maps to a UID or GID number outside this range, TrueNAS will ignore it.
Range High	Range Low and Range High set the range of UID/GID numbers the IDMap backend translates. If an external credential like a Windows SID maps to a UID or GID number outside this range, TrueNAS will ignore it.

SSSD
Compat

Generate Idmap low range based on the same algorithm that SSSD uses by default.

Kerberos

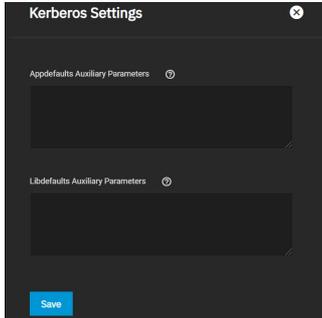
[Kerberos](#) is a web authentication protocol that uses strong cryptography to prove the identity of both client and server over an insecure network connection.

Kerberos uses “realms” and “keytabs” to authenticate clients and servers. A Kerberos realm is an authorized domain that a Kerberos server can use to authenticate a client. By default, TrueNAS creates a Kerberos realm for the local system. A [keytab](#) (“key table”) is a file that stores encryption keys for authentication.

TrueNAS SCALE allows users to configure general Kerberos settings, as well as realms and keytabs.

Kerberos Settings

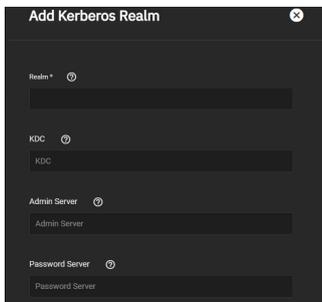
Users can configure Kerberos settings by navigating to **Directory Services** and clicking *Settings* in the *Kerberos Settings* window.



Field	Description
Appdefaults Auxiliary Parameters	Additional Kerberos application settings. See [appdefaults] in krb.conf(5) for settings and usage syntax.
Libdefaults Auxiliary Parameters	Additional Kerberos library settings. See [libdefaults] in krb.conf(5) for settings and usage syntax.

Kerberos Realms

Users can configure Kerberos realms by navigating to **Directory Services** and clicking *Add* in the *Kerberos Realms* window.



Field	Description
Realm	Enter the name of the realm.
KDC	Enter the name of the Key Distribution Center. Separate multiple values by pressing Enter.
Admin Server	Define the server that performs all database changes. Separate multiple values by pressing Enter.
Password Server	Define the server that performs all password changes. Separate multiple values by pressing Enter.

Kerberos Keytabs

Kerberos keytabs let you join an Active Directory or LDAP server without a password.

Since TrueNAS does not save the Active Directory or LDAP administrator account password in the system database, keytabs can be a security risk in some environments.

When using a keytab, create and use a less-privileged account to perform queries. TrueNAS will store that account’s password in the system database.

Create Keytab on Windows

To create the keytab on a Windows system, use the [ktpass](#) command:

```
ktpass.exe /out file.keytab /princ http/user@EXAMPLE.COM /mapuser user /ptype KRB5_NT_PRINCIPAL /crypto ALL /pass userpass
```

- file.keytab is the file to upload to the TrueNAS server.
- user is the user account name for the TrueNAS server generated in [Active Directory Users and Computers](#).
- http/user@EXAMPLE.COM is the principal name written in the format host/user.account@KERBEROS.REALM. The Kerberos realm is usually in all caps, but be sure to match the Kerberos Realm case with the realm name. See [this note](#) about using /princ for more details.
- userpass is the user’s password.
- /crypto is the cryptographic type.

Setting /crypto to ALL allows using all supported cryptographic types. You can use specific keys instead of using ALL:

- DES-CBC-CRC is backward compatible.
- DES-CBC-MD5 adheres more closely to the MIT implementation and is backward compatible.
- RC4-HMAC-NT uses 128-bit encryption.
- AES256-SHA1 uses AES256-CTS-HMAC-SHA1-96 encryption.
- AES128-SHA1 uses AES128-CTS-HMAC-SHA1-96 encryption.

Add Windows Keytab to TrueNAS

After generating the keytab, go back to **Directory Services** in TrueNAS and click *Add* in the *Kerberos Keytab* window to add it to TrueNAS.

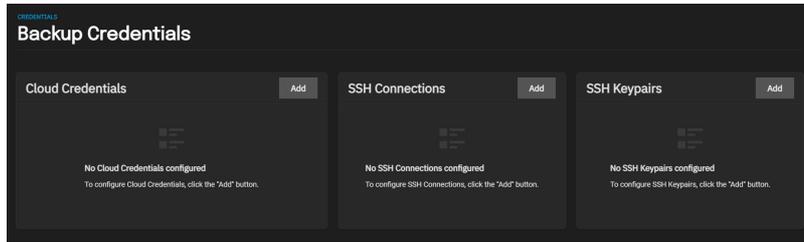
To make AD use the keytab, click *Settings* in the *Active Directory* window and select it using the *Kerberos Principal* drop-down.

When using a keytab with AD, ensure the keytab *username* and *userpass* match the *Domain Account Name* and *Domain Account Password*.

To make LDAP use a keytab principal, click *Settings* in the *LDAP* window and select the keytab using the *Kerberos Principal* drop-down.

8.4 - Backup Credentials

The **Backup Credentials** section lets users integrate TrueNAS with Cloud Storage providers and set up SSH Connections and Keypairs.



Cloud Credentials

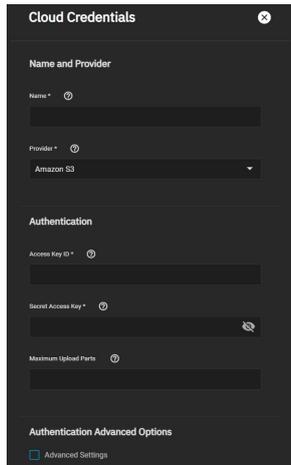
The *Cloud Credentials* window allows users to integrate TrueNAS with Cloud Storage providers.

Is this secure?

To maximize security, TrueNAS encrypts cloud credentials when saving them. However, this means that to restore any cloud credentials from a TrueNAS configuration file, you must enable *Export Password Secret Seed* when generating that [configuration backup](#). Remember to protect any downloaded TrueNAS configuration files.

We recommend users open another browser tab open and log in to the Cloud Storage Provider account you intend to link with TrueNAS. Some providers require additional information that they generate on the storage provider account page. For example, saving an Amazon S3 credential on TrueNAS could require logging in to the S3 account and generating an access key pair on the *Security Credentials > Access Keys* page.

To set up a Cloud Credential, go to **Credentials > Backup Credentials** and click *Add* in the *Cloud Credentials* window.



Enter a credential *Name* and choose a *Provider*. The rest of the options change according to the chosen *Provider*:

Amazon S3

Dolor sit, sumo unique ...

Name	Description
Access Key ID	Amazon Web Services Key ID. This is found on Amazon AWS by going through My account > Security Credentials > Access Keys (Access Key ID and Secret Access Key). Must be alphanumeric and between 5 and 20 characters.
Secret Access Key	Amazon Web Services password. If the Secret Access Key cannot be found or remembered, go to My Account > Security Credentials > Access Keys and create a new key pair. Must be alphanumeric and between 8 and 40 characters.
Maximum Upload Ports	Define the maximum number of chunks for a multipart upload. Setting a maximum is necessary if a service does not support the 10,000 chunk AWS S3 specification.

Amazon S3 Advanced Options

Name	Description
Endpoint URL	S3 API endpoint URL . When using AWS, the endpoint field can be empty to use the default endpoint for the region and automatically fetch available buckets. Refer to the AWS Documentation for a list of Simple Storage Service Website Endpoints .
Region	AWS resources in a geographic area . Leave empty to detect the bucket's correct public region. Entering a private region name allows interacting with Amazon buckets created in that region. For example, enter us-gov-east-1 to discover buckets created in the eastern AWS GovCloud region.
Disable Endpoint Region	Skip automatic detection of the Endpoint URL region. Set this when configuring a custom Endpoint URL.
User Signature Version 2	Force using Signature Version 2 to sign API requests. Set this when configuring a custom Endpoint URL.

Backblaze B2

Name	Description
Key ID	Alphanumeric Backblaze B2 Application Key ID. To generate a new application key, log in to the Backblaze account, go to the App Keys page, and add a new application key. Copy the application keyID string to this field.
Application Key	Backblaze B2 Application Key. To generate a new application key, log in to the Backblaze account, go to the App Keys page, and add a new application key. Copy the <i>applicationKey</i> string to this field.

Box

Name	Description
OAuth Client ID	The public identifier for the cloud application.
OAuth Client Secret	The secret phrase known only to the cloud application and the authorization server.
Access Token	A User Access Token for Box . An access token enables Box to verify a request belongs to an authorized session. Example token: T9cE5asGnuyYCCqIZFoWjFFHvNbvVqHjl.

DropBox

Name	Description
OAuth Client ID	The public identifier for the cloud application.
OAuth Client Secret	The secret phrase known only to the cloud application and the authorization server.
Access Token	Access Token for a Dropbox account. You must create a token from the Dropbox account before adding it here.

FTP

Name	Description
Host	FTP Host to connect. Example: ftp.example.com.
Port	FTP Port number. Leave blank to use the default port 21.
Username	A username on the FTP Host system. This user must already exist on the FTP Host.
Password	Password for the user account.

Google Cloud Storage

Name	Description
Preview JSON Service Account Key	Contents of the uploaded Service Account JSON file.
Choose File	Upload a Google Service Account credential file . The Google Cloud Platform Console creates the file.

Google Drive

Name	Description
OAuth Client ID	The public identifier for the cloud application.
OAuth Client Secret	The secret phrase known only to the cloud application and the authorization server.
Access Token	Token created with Google Drive . Access Tokens expire periodically, so you must refresh them.
Team Drive ID	Only needed when connecting to a Team Drive. The Team Drive's top-level folder ID.

Google Photos

Name	Description
OAuth Client ID	The public identifier for the cloud application.
OAuth Client Secret	The secret phrase known only to the cloud application and the authorization server.

HTTP

Name	Description
URL	HTTP host URL.

Hubic

Name	Description
Access Token	Access Token generated by a Hubic account .

Mega

Name	Description
Username	MEGA account username.
Password	MEGA account password.

Microsoft Azure Blob Storage

Name	Description
Account Name	Microsoft Azure account name.
Account Key	Base64 encoded key for Azure Account.

Microsoft One Drive

Name	Description
OAuth Client ID	The public identifier for the cloud application.
OAuth Client Secret	The secret phrase known only to the cloud application and the authorization server.
Access Token	Microsoft Onedrive Access Token . Log in to the Microsoft account to add an access token.
Drives List	Drives and IDs registered to the Microsoft account. Selecting a drive also fills the Drive ID field.
Drive Account Type	Type of Microsoft account. Logging in to a Microsoft account selects the correct account type. Options: Personal, Business, Document_Library
Drive ID	Unique drive identifier. Log in to a Microsoft account and choose a drive from the Drives List drop-down to add a valid ID.

OpenStack Swift

Name	Description
User Name	Openstack user name (OS_USERNAME) from an OpenStack credentials file .
API Key or Password	Openstack API key or password. This is the OS_PASSWORD from an OpenStack credentials file .
Authentication URL	Authentication URL for the server. This is the OS_AUTH_URL from an OpenStack credentials file .
Auth Version	AuthVersion - optional - set to (1,2,3) if your auth URL has no version (rclone documentation).

Advanced Options

Name	Description
Tenant Name	This is the OS_TENANT_NAME from an OpenStack credentials file .
Tenant ID	Tenant ID - optional for v1 auth, this or tenant required otherwise (rclone documentation).
Auth Token	Auth Token from alternate authentication - optional (rclone documentation).

Endpoint Advanced Options

Name	Description
Region Name	Region name - optional (rclone documentation).

Storage URL	Storage URL - optional (rclone documentation).
Endpoint Type	Endpoint type to choose from the service catalogue. Public is recommended, see the rclone documentation .

pCloud	
Name	Description
OAuth Client ID	The public identifier for the cloud application.
OAuth Client Secret	The secret phrase known only to the cloud application and the authorization server.
Access Token	pCloud Access Token . These tokens can expire and require an extension.
Hostname	Enter the hostname to connect to.

SFTP	
Name	Description
Host	SSH Host to connect to.
Port	SSH port number. Leave empty to use the default port 22.
Username	SSH Username.
Password	Password for the SSH Username account.
Private Key ID	Import the private key from an existing SSH keypair or select Generate New to create a new SSH key for this credential.

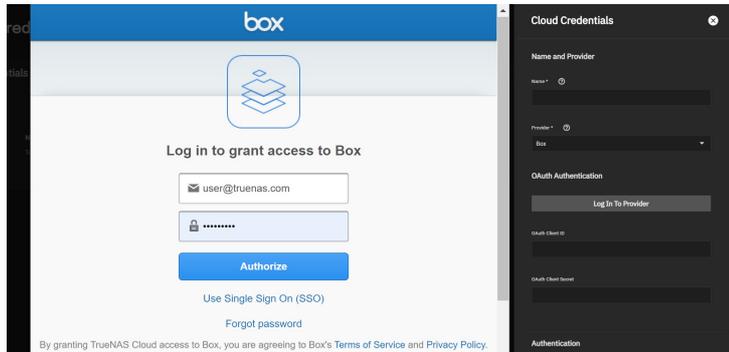
WebDav	
Name	Description
URL	URL of the HTTP host to connect to.
WebDav Service	Name of the WebDAV site, service, or software being used.
Username	WebDAV account username.
Password	WebDAV account password.

Yandex	
Name	Description
OAuth Client ID	The public identifier for the cloud application.
OAuth Client Secret	The secret phrase known only to the cloud application and the authorization server.
Access Token	Yandex Access Token .

Enter the required *Authentication* strings to enable saving the credential.

Automatic Authentication

Some providers can automatically populate the required *Authentication* strings by logging in to the account. To automatically configure the credential, click *Login to Provider* and entering your account username and password.



We recommend verifying the credential before saving it.

SSH Connections

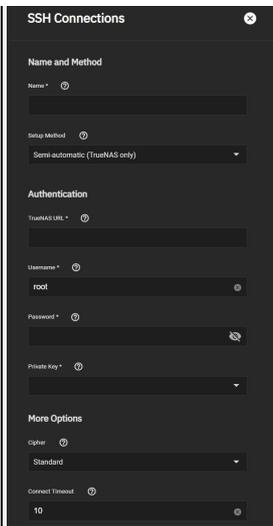
The *SSH Connections* window in the **Backup Credentials** screen allows users establish [Secure Socket Shell \(SSH\)](#) connections.

To begin setting up a SSH Connection, navigate to **Credentials > Backup Credentials** and click the *Add* button in the *SSH Connections* window.

Create a Connection

Semi-Automatic

Semi-automatic simplifies setting up an SSH connection with another FreeNAS or TrueNAS system without logging in to that system to transfer SSH keys. This requires an SSH keypair on the local system and administrator account credentials for the remote TrueNAS. You must configure the remote system to allow root access with SSH. You can generate the keypair as part of the semiautomatic configuration or a manually created one in **Backup Credentials**.



Name and Method

Name	Description
Name	Name of this SSH connection. SSH connection names must be unique.
Setup Method	<i>Manual</i> requires configuring authentication on the remote system. This can include copying SSH keys and modifying the root user account on that system. <i>Semi-automatic</i> only works when configuring an SSH connection with a remote TrueNAS system. This method uses the URL and login credentials of the remote system to connect and exchange SSH keys.

Authentication

Name	Description
TrueNAS URL	Hostname or IP address of the remote system. A valid URL scheme is required. Example: <code>https://10.231.3.76</code>
Username	Username for logging in to the remote system.
Password	User account password for logging into the remote system.
Private Key	Choose a saved SSH Keypair or select Generate New to create a new keypair and use it for this connection.

More Options

Name	Description
Cipher	<i>Standard</i> is most secure, but has the greatest impact on connection speed. <i>Fast</i> is less secure than Standard but can give reasonable transfer rates for devices with limited cryptographic speed. <i>Disabled</i> removes all security in favor of maximizing connection speed. Disabling the security should only be used within a secure, trusted network.
Connect Timeout	Time (in seconds) before the system stops attempting to establish a connection with the remote system.

Be sure to use a valid URL scheme for the remote TrueNAS URL. Leave the username as `root` and enter the account password for the remote TrueNAS system. You can import the private key from a previously created SSH keypair or create one with a new SSH keypair.

Saving the new configuration automatically opens a connection to the remote TrueNAS and exchanges SSH keys.

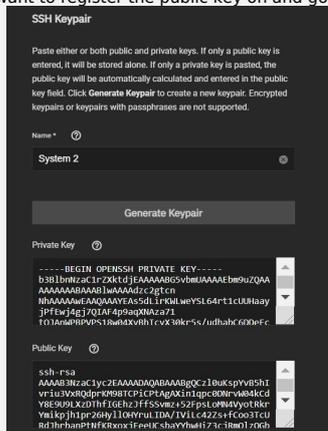
Manual

To manually set up an SSH connection, you must copy a public encryption key from the local system to the remote system. A manual setup allows a secure connection without a password prompt.

Adding a Public SSH Key to the TrueNAS Root Account

Log in to the TrueNAS system that generated the SSH keypair and go to **Credentials > Backup Credentials**. Click the `+` icon to add a new credential. Open the keypair for the SSH connection and copy the text of the public SSH key or download the public key as a text file.

Log in to the TrueNAS system you want to register the public key on and go to **Credentials > Local Users**. Edit the `root` account. Paste the SSH public key



text into the **SSH Public Key** field.

Start by generating a new SSH keypair in **Credentials > Backup Credentials**. Copy or download the value for the public key. Add the public key to the remote NAS. If the remote NAS is not a TrueNAS system, please see the documentation for that system for instructions on adding a public SSH key.

Manually Configuring the SSH Connection on the Local TrueNAS

Log back in to the local TrueNAS system. Go to **Credentials > Backup Credentials** and add a new SSH connection. Change the setup method to *Manual*.

Name and Method

Name	Description
Name	Name of this SSH connection. SSH connection names must be unique.
Setup Method	<i>Manual</i> requires configuring authentication on the remote system. This can include copying SSH keys and modifying the root user account on that system. <i>Semi-automatic</i> only works when configuring an SSH connection with a remote TrueNAS system. This method uses the URL and login credentials of the remote system to connect and exchange SSH keys.

Authentication

Name	Description
Host	Hostname or IP address of the remote system. A valid URL scheme is required. Example: <code>https://10.231.3.76</code>
Port	Port number on the remote system to use for the SSH connection.
Username	Username for logging in to the remote system.
Private Key	Choose a saved SSH Keypair or select <i>Generate New</i> to create a new keypair and use it for this connection.
Remote Host Key	Remote system SSH key for this system to authenticate the connection. When all other fields are properly configured, click <i>DISCOVER REMOTE HOST KEY</i> to query the remote system and automatically populate this field.

Discover Remote Host Key connects to the remote host and attempts to copy the key string to the related TrueNAS field.

More Options

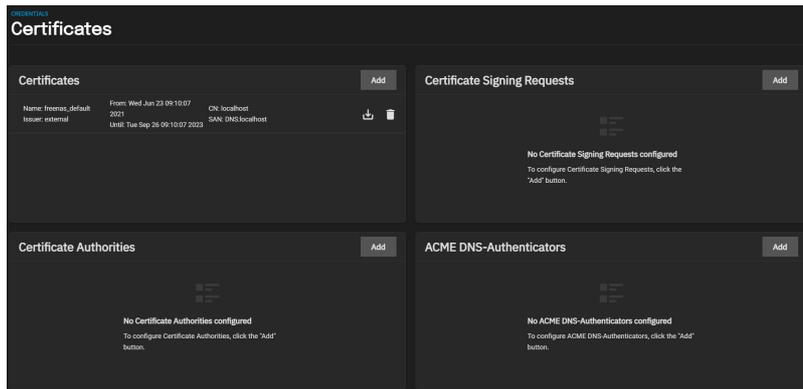
Name	Description
Cipher	<i>Standard</i> is most secure, but has the greatest impact on connection speed. <i>Fast</i> is less secure than Standard but can give reasonable transfer rates for devices with limited cryptographic speed. <i>Disabled</i> removes all security in favor of maximizing connection speed. Disabling the security should only be used within a secure, trusted network.
Connect Timeout	Time (in seconds) before the system stops attempting to establish a connection with the remote system.

Select the private key from the SSH keypair that you used to transfer the public key on the remote NAS.

8.5 - Certificates

- - [Certificates](#)
 - [Certificate Signing Requests](#)
 - [Certificate Authorities](#)
 - [ACME DNS-Authenticators](#)

The Certificates section contains all the information for certificates, certificate signing requests, certificate authorities, and DNS-authenticators. TrueNAS comes equipped with an internal, self-signed certificate that enables encrypted access to the web interface, but users can make custom certificates for authentication and validation while sharing data.



Certificates

By default, TrueNAS comes equipped with an internal, self-signed certificate that enables encrypted access to the web interface, but users can import and create more certificates by clicking **Add** in the **Certificates** window.

Identifier and Type

The **Identifier and Type** step lets users name the certificate and choose whether to use it for internal or local systems, or import an existing certificate. Users can also select a predefined certificate extension from the **Profiles** drop-down list.

Certificate Options

The **Certificate Options** step provides options for choosing the Signing Certificate Authority (CSR), what type of private key type to use (as well as the number of bits in the key used by the cryptographic algorithm), the cryptographic algorithm the certificate uses, and how many days the certificate authority lasts.

Certificate Subject

The **Certificate Subject** step lets users define the location, name, and email for the organization using the certificate. Users can also enter the system [fully-qualified hostname \(FQDN\)](#) and any additional domains for multi-domain support.

Extra Constraints

The **Extra Constraints** step contains certificate extension options.

- **Basic Constraints:** Enable to limit the path length for a certificate chain.
- **Authority Key Identifier:** Enable to provide a means of identifying the public key corresponding to the private key used to sign a certificate.
- **Key Usage:** Enable to define the purpose of the public key contained in a certificate.
- **Extended Key Usage:** Enable to further refine key usage extensions.

Certificate Signing Requests

The **Certificate Signing Requests** section allows users configure the message(s) the system sends to a registration authority of the public key infrastructure to apply for a digital identity certificate.

Identifier and Type

The **Identifier and Type** step lets users name the certificate signing request (CSR) and choose whether to create a new CSR or import an existing CSR. Users can also select a predefined certificate extension from the **Profiles** drop-down list.

Certificate Options

The **Certificate Options** step provides options for choosing what type of private key type to use, the number of bits in the key used by the cryptographic algorithm, and the cryptographic algorithm the CSR uses.

Extra Constraints

The **Extra Constraints** step contains certificate extension options.

- **Basic Constraints:** Enable to limit the path length for a certificate chain.
- **Authority Key Identifier:** Enable to provide a means of identifying the public key corresponding to the private key used to sign a certificate.
- **Key Usage:** Enable to define the purpose of the public key contained in a certificate.
- **Extended Key Usage:** Enable to further refine key usage extensions.

Certificate Authorities

The Certificate Authorities section lets users set up a certificate authority (CA) that certifies the ownership of a public key by the named subject of the certificate.

Identifier and Type

The **Identifier and Type** step lets users name the CA and choose whether to create a new CA or import an existing CA. Users can also select a predefined certificate extension from the **Profiles** drop-down list.

Certificate Options

The **Certificate Options** step provides options for choosing what type of private key to use (as well as the number of bits in the key used by the cryptographic algorithm), the cryptographic algorithm the CA uses, and how many days the CA lasts.

Certificate Subject

The **Certificate Subject** step lets users define the location, name, and email for the organization using the certificate. Users can also enter the system [fully-qualified hostname \(FQDN\)](#) and any additional domains for multi-domain support.

Extra Constraints

The **Extra Constraints** step contains certificate extension options.

- **Basic Constraints:** Enable to limit the path length for a certificate chain.
- **Authority Key Identifier:** Enable to provide a means of identifying the public key corresponding to the private key used to sign a certificate.
- **Key Usage:** Enable to define the purpose of the public key contained in a certificate.
- **Extended Key Usage:** Enable to further refine key usage extensions.

ACME DNS-Authenticators

The Automatic Certificate Management Environment (ACME) DNS-Authenticators screen allows users to automate certificate issuing and renewal. The user must verify ownership of the domain before certificate automation is allowed.

The system requires an ACME DNS Authenticator and CSR to configure ACME certificate automation.

Users must name the authenticator and choose a DNS provider and configure any required authenticator attributes.

If you select [Cloudflare](#) as the authenticator, you must enter your Cloudflare account email address, API Key, and API Token.

If you select [Route53](#) as the authenticator, you must enter you Route53 Access Key ID and Secret Access Key.

8.6 - 2FA (Two-Factor Authentication)

- [2FA Options](#)
- [Enabling Two-Factor Authentication.](#)
- [Using 2FA to Log in to TrueNAS](#)

Two-factor authentication (2FA) is great for increasing security.

TrueNAS offers 2FA to ensure that entities cannot use a compromised administrator *root* password to access the administrator interface.

You need a mobile device with the current time and date that has Google Authenticator installed to use 2FA.

Two-Factor authentication is time based and requires that the system is set correctly. Making sure NTP is functional before enabling is strongly recommended!

What is 2FA and why should I enable it?

2FA adds an extra layer of security to your system to prevent someone from logging in, even if they have your password. 2FA requires you to verify your identity using a randomized 6-digit code that regenerates every 30 seconds (unless modified) to use when you log in.

Benefits

Unauthorized users can't log in since they won't have the randomized 6-digit code.

Authorized employees can securely access systems from any device or location without jeopardizing sensitive information.

Internet access on the TrueNAS system is not required to use 2FA.

Drawbacks

Requires an app to generate 2FA code.

If the 2FA code isn't working or users can't get it, the system is inaccessible through the UI and SSH (if enabled).

If the device with the 2FA app isn't available, you can use the system CLI to bypass 2FA with administrative IPMI or by physically accessing the system.

To unlock 2FA in the CLI, enter: `midclt call auth.twofactor.update '{ "enabled":false }'`

2FA Options



Two-factor authentication is time-based and requires that you correctly set the system time.

User Settings

Name	Description
One Time Password (OTP) Digits	The number of digits in the One-Time Password. The default is 6, which is Google's standard OTP length. Check your app/device settings before selecting this.
Interval	The lifespan (in seconds) of each OTP. Default is 30 seconds. The minimum is 5 seconds.
Window	Extends password validity beyond the <i>Interval</i> setting. For example, 1 means that one password before and after the current one is valid, leaving three valid passwords. Extending the window is useful in high-latency situations.
Enable Two-Factor Auth for SSH	Enable 2FA for system SSH access. We recommend leaving this DISABLED until after you successfully test 2FA with the UI.

System Generated Settings

Name	Description
Secret (Read-only)	The secret TrueNAS creates and uses to generate OTPs when you first enable 2FA.
Provisioning URI (includes Secret - Read-only)	The URI used to provision an OTP. TrueNAS encodes the URI (which contains the secret) in a QR Code. To set up an OTP app like Google Authenticator, use the app to scan the QR code or enter the secret manually into the app. TrueNAS produces the URI when you first activate 2FA.

Enabling Two-Factor Authentication.

Video Tutorial

This short video demonstrates adding 2FA.



Video URL: <https://www.truenas.com/docs/files/scaleangelfish2fa.mp4>

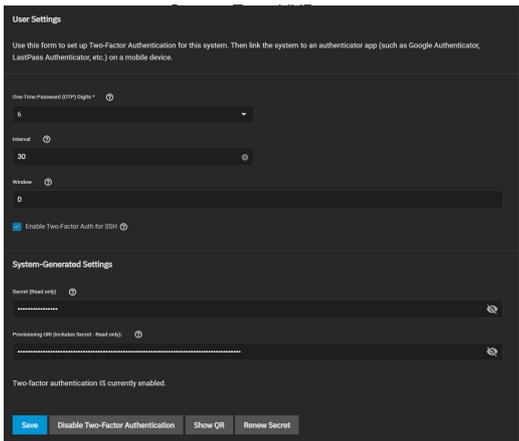
Play Video

Play

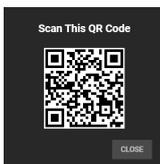
Set up a second 2FA device as a backup before proceeding.

1 Go to **Credentials** > **2FA**.

2 Click **Enable Two Factor Authentication**, then click **Confirm**.



3 Click **Show QR**.



Fullscreen

This is a modal window.

Beginning of dialog window. Escape will cancel and close the window.

Text

Color White Transparency Opaque

4 Start Google Authenticator on the mobile device and scan the QR code.

Color Black Transparency Opaque

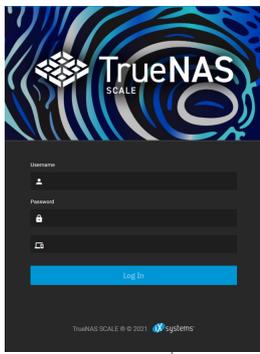
Using 2FA to Log in to TrueNAS

Enabling 2FA changes the login process for both the TrueNAS web interface and SSH logins:

Web Interface

The login screen adds another field for the randomized authenticator code. If this field isn't immediately visible, try refreshing the browser.

Enter the code from the mobile device (without the space) in the login window with the *root Username* and *Password*.



Style

ily

Sans-Serif

all settings to the default values Done

alog

g window.

SSH Logins

1 Confirm that you set **Enable Two-Factor Auth for SSH in Credentials > 2FA**. 2 Go to **System Settings > Services** and edit the **SSH** service. Set **Log in as Root with Password**, then click **Save**. Toggle the **SSH** service and wait for the status to show that it is running. 3 Open the Google Authentication app on your mobile device. 4 Open a terminal and SSH into the system using its hostname or IP address, *root* account username and password, and the 2FA code.

```
PS C:\Users\TrueNAS User> ssh root@10.10.10.10
Password:
One-time password (OTP) for "root":
Linux truenas.local 5.10.79+truenas #1 SMP Mon Oct 11 16:48:36 UTC 2021 x86_64

      TrueNAS (c) 2009-2021, iXsystems, Inc.
      All rights reserved.
      TrueNAS code is released under the modified BSD license with some
      files copyrighted by (c) iXsystems, Inc.

      For more information, documentation, help or support, go here:
      http://truenas.com

Welcome to TrueNAS
last login: Tue Oct 12 10:16:25 2021

Warning: settings changed through the CLI are not written to
the configuration database and will be reset on reboot.

root@truenas[-]#
```

9 - Virtualization

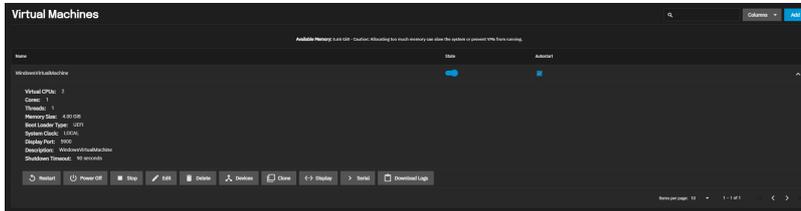
The Virtualization section allows users to set up Virtual Machines (VMs) to run alongside TrueNAS. Delegating processes to VMs reduces the load on the physical system, which means users can utilize additional hardware resources. Users can customize six different segments of a VM when creating one in TrueNAS SCALE.

What system resources do VMs require?

TrueNAS assigns a portion of system RAM and a new zvol to each VM. While a VM is running, these resources are not available to the host computer or other VMs.

TrueNAS VMs use the [KVM](#) virtual machine software. This type of virtualization requires an x86 machine running a recent Linux kernel on an Intel processor with VT (virtualization technology) extensions or an AMD processor with SVM extensions (also called AMD-V). Users cannot create VMs unless the host system supports these features.

To verify that you have Intel VT or AMD-V, open the **Shell** and run `egrep '^flags.*(vmx|svm)' /proc/cpuinfo`. If device information appears, your system has VT. You can also check the processor model name (in `/proc/cpuinfo`) on the vendor's website.



1 Operating System

The **Operating System** menu lets users choose the VM operating system type, the time it uses, its boot method, and its display type.

The menu also lets users set a shutdown timeout duration and IP address type, as well as set whether the VM should start when the system boots or have a display.

2 CPU and Memory

The **CPU and Memory** menu lets users select how many virtual CPUs to allocate to the virtual machine, how many cores per virtual CPU socket, and how many threads per core.

This menu also has options for CPU mode and model, and how much RAM to allocate for the VM.

3 Disks

The **Disks** menu lets users choose to either create a new zvol on an existing dataset for a disk image or use an existing zvol or file for the VM.

Users may also select the disk type, zvol location, and how much space to allocate to the zvol.

4 Network Interface

The **Network Interface** menu provides options for the adapter type, Mac address, and which physical interface to associate with the VM.

5 Installation Media

The **Installation Media** menu lets users decide if they want to choose an installation media image on a dataset or upload one from the local machine.

6 GPU

The **GPU** menu allows users to select a graphics processing unit (GPU) for the VM. It also provides the option to hide the VM from the Microsoft Reserved Partition (MSR) on Windows systems.

Ready to get started? Choose a topic or article from the left-side Navigation pane. Click the < symbol to expand the menu to show the topics under this section.

9.1 - Creating and Managing VMs

- [Creating a Virtual Machine](#)
 - [Adding and Removing Devices](#)
- [Managing the Virtual Machine](#)
- [Installing an OS](#)

A Virtual Machine (VM) is an environment on a host computer that you can use as if it were a separate, physical computer. Users can use VMs to run multiple operating systems simultaneously on a single computer. Operating systems running inside a VM see emulated virtual hardware rather than the host computer physical hardware. VMs provide more isolation than jails but also consumes more system resources.

What system resources do VMs require?

TrueNAS assigns a portion of system RAM and a new zvol to each VM. While a VM is running, these resources are not available to the host computer or other VMs.

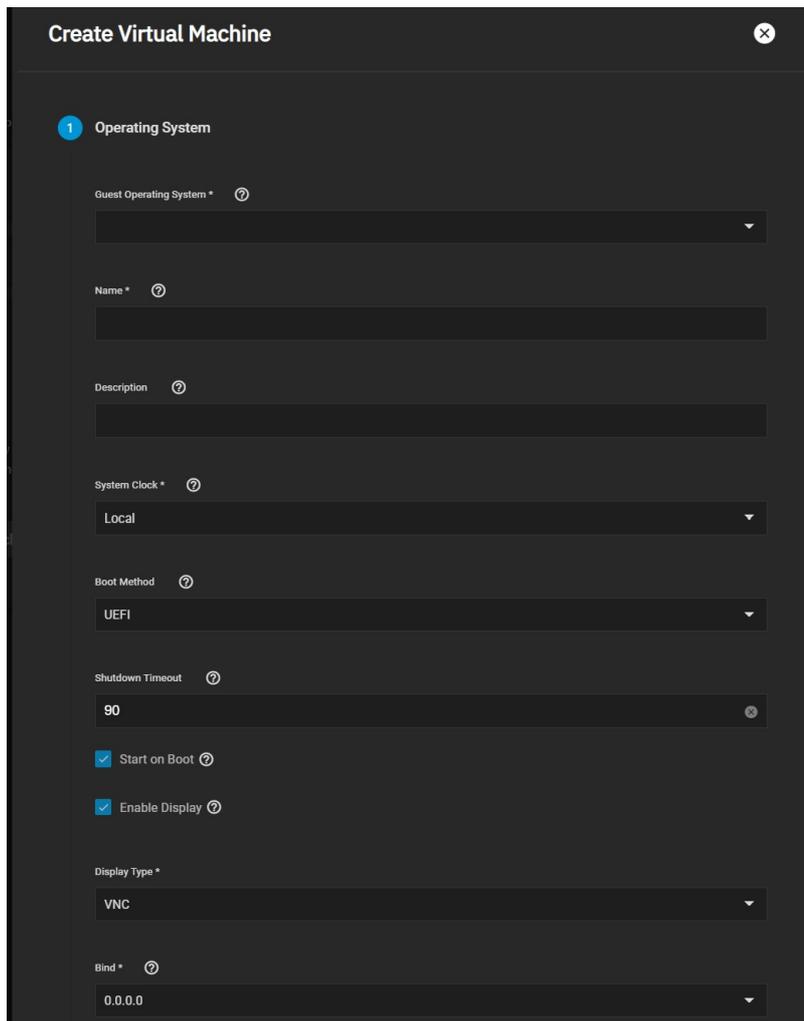
TrueNAS VMs use the [KVM](#) virtual machine software. This type of virtualization requires an x86 machine running a recent Linux kernel on an Intel processor with VT (virtualization technology) extensions or an AMD processor with SVM extensions (also called AMD-V). Users cannot create VMs unless the host system supports these features.

To verify that you have Intel VT or AMD-V, open the **Shell** and run `egrep '^flags.*(vmx|svm)' /proc/cpuinfo`. If device information appears, your system has VT. You can also check the processor model name (in `/proc/cpuinfo`) on the vendor's website.

Creating a Virtual Machine

Before creating a virtual machine, you need an installer .iso or image file for the OS you intend to install, and a [storage pool](#) available for both the virtual disk and OS install file.

To create a new VM, go to **Virtualization** and click **Add** (or **Add Virtual Machines**). Configure each category of the VM according to your specifications, starting with the **Operating System**.



Specific Options

Operating System

Field	Description
Guest Operating System	Choose the VM operating system type.
Name	Enter an alphanumeric name for the virtual machine.
Description	Enter a description (optional).
System Clock *	The VM system time. Default is Local.
Boot Method	Select UEFI for newer operating systems or Legacy BIOS for older operating systems that only support BIOS booting.
Shutdown Timeout	The time in seconds the system waits for the VM to cleanly shut down. During system shutdown, the system initiates poweroff for the VM after the shutdown timeout has expired.
Start on Boot	Set to start this VM when the system boots.
Enable Display	Enable a Display (Virtual Network Computing) remote connection. Requires UEFI booting.
Display Type	Select either VNC or SPICE
Bind	Display network interface IP address. The primary interface IP address is the default. A different interface IP address can be chosen.

CPU and Memory

Field	Description
Virtual CPUs	The number of virtual CPUs to allocate to the virtual machine. The maximum is 16, or fewer if the host CPU limits the maximum. The VM operating system might also have operational or licensing restrictions on the number of CPUs.
Cores	Specify the number of cores per virtual CPU socket. The product of vCPUs, cores, and threads must not exceed 16.
Threads	Specify the number of threads per core. The product of vCPUs, cores, and threads must not exceed 16.
CPU Mode	Choose between Custom, Model, and Passthrough modes.
CPU Model	Select a CPU model to emulate.
Memory Size	Allocate RAM for the VM. Minimum value is 256 MiB. This field accepts human-readable input (Ex. 50 GiB, 500M, 2 TB). If units are not specified, the value defaults to bytes.

Disks

Field	Description
Create new disk image	Create a new zvol on an existing dataset to use as a virtual hard drive for the VM.
Use existing disk image	Use an existing zvol or file for the VM.
Select Disk Type	Select desired disk type (AHCI or VirtIO).
Zvol Location	Select a dataset for the new zvol.
Size	Allocate space for the new zvol. (Examples: 500 KIB, 500M, 2 TB) MiB. Units smaller than MiB are not allowed.

Network Interface

Field	Description
Adapter Type	Intel e82545 (e1000) emulates the same Intel Ethernet card and provides compatibility with most operating systems. VirtIO provides better performance when the operating system installed in the VM supports VirtIO paravirtualized network drivers.
Mac Address	Enter the desired address into the field to override the randomized MAC address.
Attach NIC	Select the physical interface to associate with the VM.

Installation Media

Field	Description
Choose Installation Media Image	Browse to the operating system installer image file.
Upload an Installer Image File	Set to display image upload options.

GPU

Field	Description
Hide from MSR	Enable to hide the GPU from the Microsoft Reserved Partition (MSR)
GPU's	Select a physical GPU on your system to use for the VM.

Additional notes:

Compare the recommended specifications for your guest operating system with the available host system resources when allocating virtual CPUs, cores, threads, and memory size.

Don't allocate too much memory to a VM.

Activating a VM with all available memory allocated to it can slow the host system or prevent other VMs from starting.

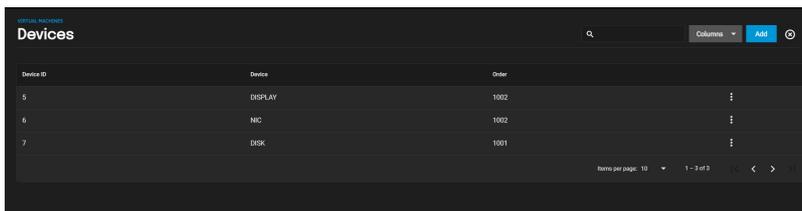
We recommend using the **AHCI Disk Type** for Windows VMs.

The **VirtIO** network interface requires a guest OS that supports VirtIO paravirtualized network drivers.

Adding and Removing Devices

After creating the VM, add and remove virtual devices by expanding the VM entry in **Virtualization** and clicking

Devices.

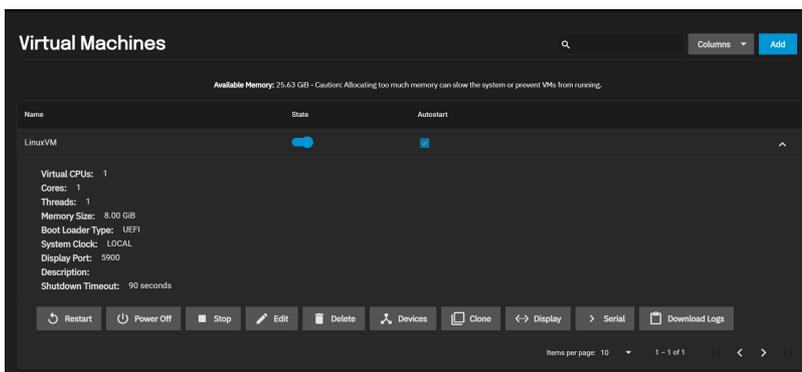


Device notes:

- Virtual machine attempt to boot from devices according to the **Device Order**, starting with **1000**, then ascending.
- **CD-ROM** devices allow booting a VM from a CD-ROM image like an installation CD. The CD image must be available in the system storage.

Managing the Virtual Machine

After creating the VM and configuring any devices for it, manage the VM by expanding its entry in **Virtualization**.



When the VM is active, it displays options for

Display and

Serial connections.

If the display connection screen appears distorted, try adjusting the display device resolution.

Use the **State** toggle or click

Stop to follow a standard shut down procedure to do a clean shut down of the running VM. Click **Power Off** halts and deactivates the VM, similar to unplugging a computer.

If the VM you created has no Guest OS installed, The VM **State** toggle and **Stop** button might not function as expected. The **State** toggle and **Stop** button send an ACPI power down command to the VM operating system, but since an OS is not installed, the commands time out. Use the **Power Off** button instead.

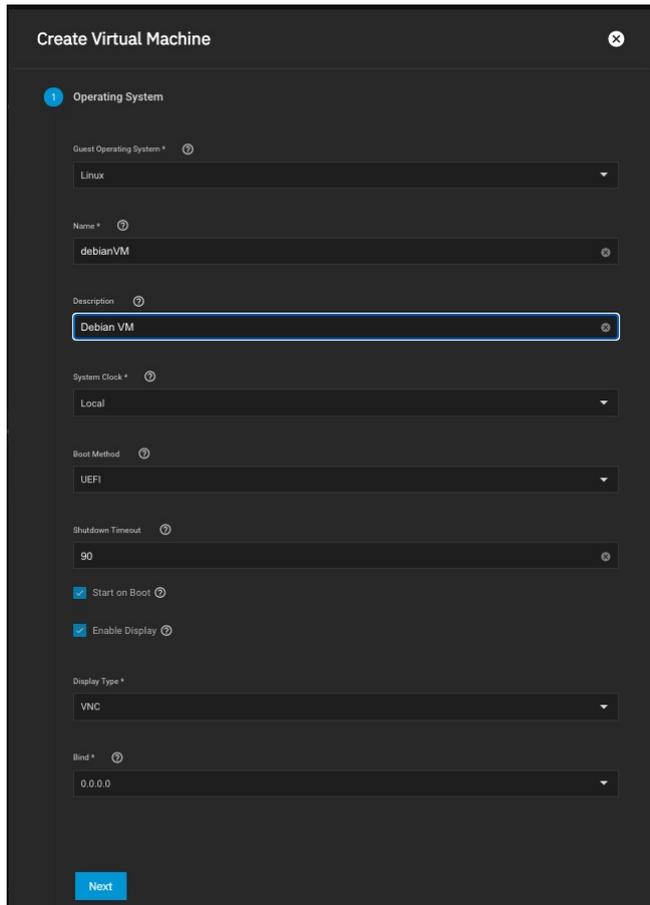
Installing an OS

When the VM is configured in TrueNAS and has an OS .iso file attached, you can start the VM and begin installing the operating system.

Some operating systems can require specific settings to function properly in a virtual machine. For example, vanilla Debian can require advanced partitioning when installing the OS. Refer to the documentation for your chosen operating system for tips and configuration instructions.

Here is an example of installing a Debian OS in a TrueNAS VM. The Debian .iso was uploaded to the TrueNAS system and attached to the VM as a CD-ROM device.

- Click on the **Virtualization** menu then click **ADD** to start the VM creation process using the wizard.



VM values entered for the Debian Example.

Operating System:

- Guest Operating System: Linux
- Name: debianVM
- Description: Debian VM

CPU and Memory:

- Change the memory size to 1024 MiB.

Disks:

- Select **Create new disk image**.
- Select the Zvol Location.
- Change the size to 30 GiB.

Network Interface:

- Attach NIC: Select the physical interface to associate with the VM.

Installation Media:

- In this case the installation ISO is uploaded to `/mnt/tank2/isostorage/`. Click on the installation ISO, `debian-11.0.0-amd64-netinst.iso`.
- If the ISO is or was not uploaded, you need to set **Upload an installer image file**, select a dataset to store the ISO, click **Choose file**, then click **Upload**. Wait for the upload to complete (this can take some time).

GPU:

- Leave the default values.

Confirm Options

- Verify the information is correct and then click **Submit**.

- After the VM is created, start it by expanding the VM entry (select the down-pointing arrow to the right of the VM name) and click **Start**.
- Click **Display** to open a virtual monitor to the VM and see the Debian Graphical Installation screens.

Debian Install Example.

Debian Graphical Install

- Press **Return** to start the Debian Graphical Install.
- Language: English
- Location: United States
- Keymap: American English

Installation begins

- Continue if the network configuration fails.
- Do not configure the network at this time.
- Enter a name in **Hostname**.
- Enter the root password and re-enter the root password.
- Enter a name in **New User**.
- Select the username for your account (it should already be filled in).
- Enter and re-enter the password for the user account.
- Choose the time zone, *Eastern* in this case.

Disk Detection should begin

- Partition disks: select **Guided - use entire disk**.
- Select the available disk.
- Select **All files in one partition** (recommended for new users).
- Select **Finish partitioning and write changes to disk**.
- Select **Yes** to **Write the changes to disks?**

Installing the base system

- Select **No** to the question **Scan extra installation media**.
- Select **Yes** when asked **Continue without a network mirror**.

Installing Software

- Select **No** when asked **Participate in the package usage survey**.
- Select **Standard** system utilities.
- Click **Continue** when the installation finishes.

- After the Debian installation finishes, close the display window.
- In the VM's expanded section click **Power Off** to stop the new VM.
- Click **Devices**.
- Remove the CD-ROM from the devices by clicking the **!** and selecting **Delete**. Click **Delete Device**.
- Click **Virtualization** and expand the new VM by selecting the down-pointing arrow to the right.
- Click **Start**.
- Click **Display**.

The grub file does not run when you start the VM. This can be done manually after each start. At the shell prompt:

- Type **FS0:** Return.
- Type **cd EFI** Return.
- Type **cd Debian** Return.
- Type **grubx64.efi** Return.

To insure it starts automatically, you need to create the startup.nsh file at the root directory on the vm. At the shell prompt type **edit startup.nsh**. In the editor type:

- Type **FS0:** Return.
- Type **cd EFI** Return.
- Type **cd Debian** Return.
- Type **grubx64.efi** Return.
- Type the **Control+s** keys(Command+s for Mac OS) Return.
- Type the **Control+q** keys to quit.

Close the display window

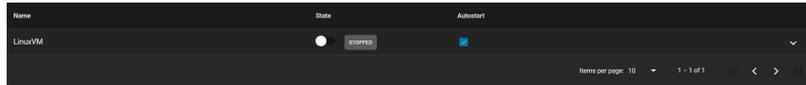
To test if it now boots up on startup:

- Power off the VM.
- Click **Start**.
- Click **Display**.
- Log into your Debian VM.

9.2 - Accessing NAS From a VM

If you want to access your TrueNAS SCALE directories from a VM, you must create a bridge interface for the VM to use.

Go to **Virtualization**, find the VM you want to use to access TrueNAS storage, and toggle it off.



Go to **Network** and find the active interface you used as the VM parent interface. Note the interface IP Address and subnet mask.

You can also get the IP address and subnet mask by going to **Shell** and entering `ip a`.

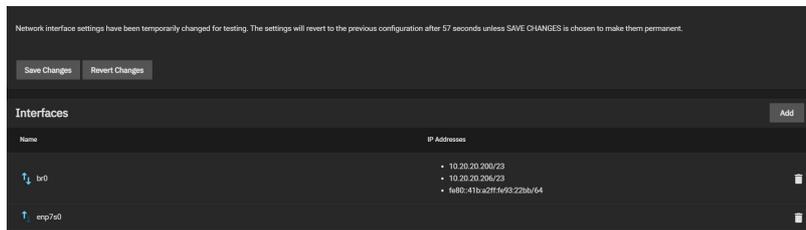
```
root@trueNAS:~# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host loop
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: enp7s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 08:00:27:4d:5d:c7:bd brd ff:ff:ff:ff:ff:ff
    inet 10.20.20.200/23 brd 10.20.21.255 scope global dynamic enp7s0
```

Click the interface. Uncheck the **DHCP** box, then click **Apply**.

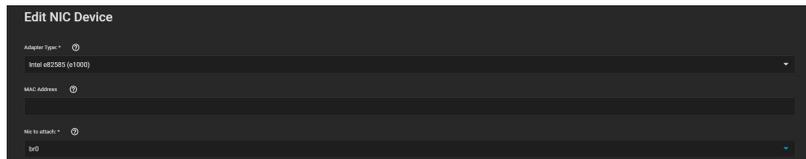


Click **Add** in the **Interfaces** window. Select **Bridge** for the **Type** and give it a name (must be in `brX` format). Check the **DHCP** box, then select the active interface on the **Bridge Members** drop-down list. Click **Add** under **IP Addresses** and enter the active interface's IP and subnet mask.

Click **Apply**, then click **Test Changes**. Once TrueNAS finishes testing the interface, click **Save Changes**.



Go to **Virtualization**, expand the VM you want to use to access TrueNAS storage, and click **Devices**. Click **Edit**. Select the new bridge interface from the **Nic to attach:** drop-down list, then click **Save**.



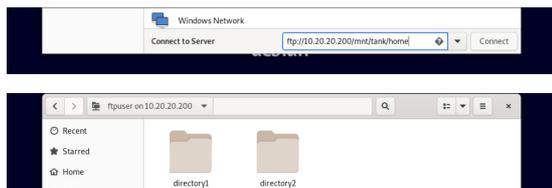
You can now access your TrueNAS storage from the VM. You might have to set up [shares](#) or [users](#) with home directories to access certain files.

Examples

Linux

Linux VMs can access TrueNAS storage using FTP, SMB, and NFS.

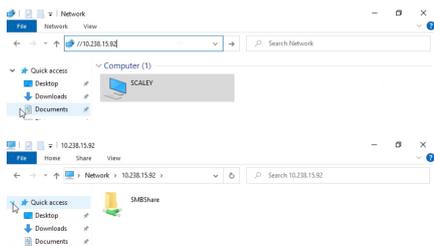
In the example below, the Linux VM is using FTP to access a user's home directory on TrueNAS.



Windows

Windows VMs can access TrueNAS storage using FTP and SMB.

In the example below, the Windows VM accessing an SMB share on TrueNAS.

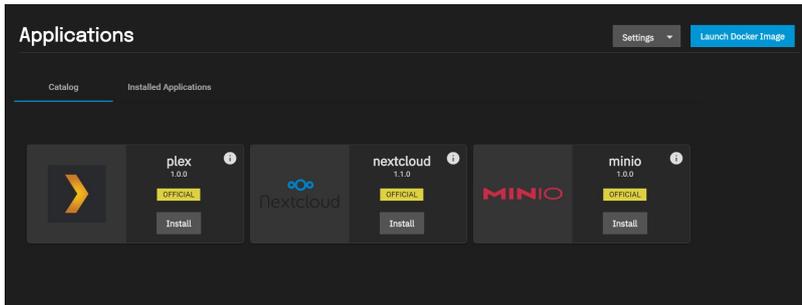


10 - Apps

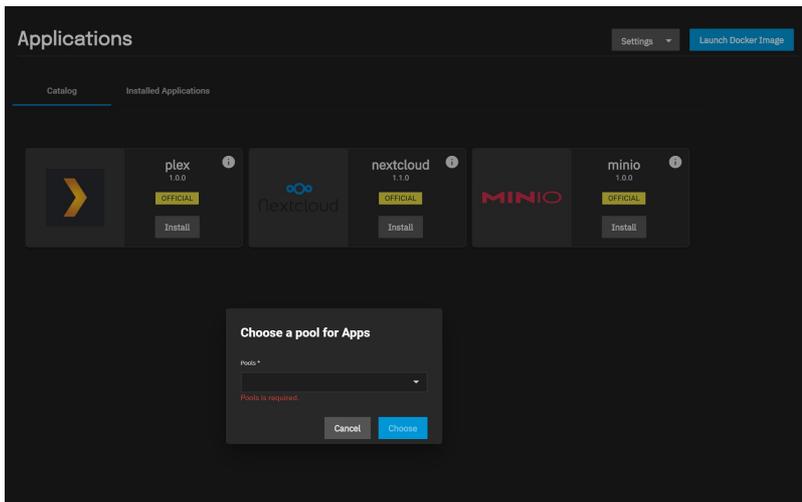
10.1 - Using SCALE Apps

- [Official Applications](#)
- [Custom Applications](#)
 - [Volume\(s\) Access](#)
- [Deploying the Application](#)
- [Accessing the Shell in an Active Container](#)

Both pre-built official containers and custom application containers can be deployed using the *Apps* page in the Scale web interface.

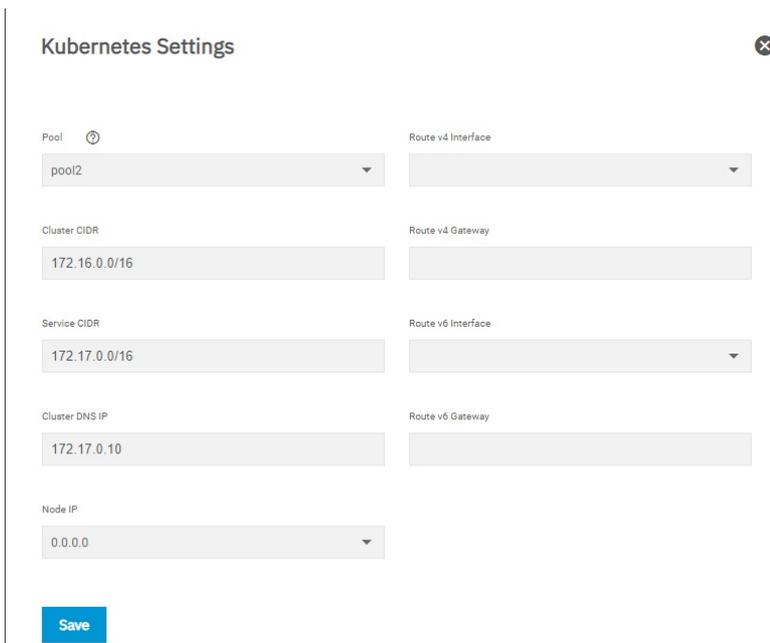


The UI will ask to use a storage pool for Applications.



We recommend users keep the container use case in mind when choosing a pool. Select a pool that has enough space for all the application containers you intend to use. TrueNAS creates an *ix-applications* dataset on the chosen pool and uses it to store all container-related data.

You can find additional options for configuring general network interfaces and IP addresses for application containers in **Apps > Settings > Advanced Settings**.



Official Applications

Official containers are pre-configured to only require a name during deployment.

Application Name *

Image

Settings

Image Repository * ?

Plex Claim Token

Image Tag ?

Advertise IP

Image Pull Policy ?

Timezone

Extra Environment Variables

Environment Variable Name	Environment Variable Value
<input type="text"/>	<input type="text"/>

Plex Node Port

Host Network

Transcode Hostpath Enabled

Data Hostpath Enabled

Config Hostpath Enabled

Save

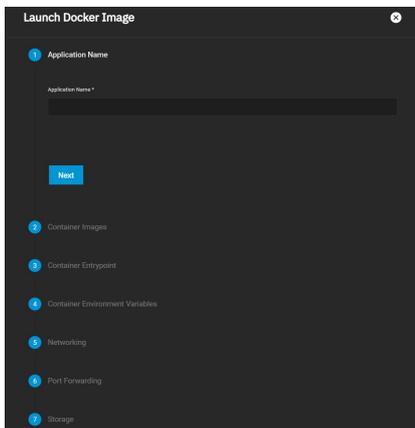
A button to open the application web interface will appear when the container is deployed and active.



Users can adjust the container settings by editing a deployed official container. Saving any changes redeploys the container.

Custom Applications

To deploy a custom application container in the Scale web interface, go to **Apps** and click *Launch Docker Image*.



TrueNAS has numerous options for custom containers that are broken down into smaller sections. These options are derived from the [Kubernetes container options](#).

Image and Policies

You will need to name the custom application and provide the online storage location (repository) that the system will use to download the container. The remaining options allow setting the image tag, defining when the image is pulled from the remote repository, how the container is updated, and defining when a container will automatically restart.

Container Settings

Define any [commands and arguments](#) to use for the image. These can override any existing commands stored in the image.

You can also [define additional environment variables](#) for the container. Some Docker images can require additional environment variables. Be sure to check the documentation for the image you're trying to deploy and add any required variables here.

Networking

To use the system IP address for the container, set *Host Networking*. The container will not be given a separate IP address and the container port number will be appended to the end of the system IP address. See the [Docker documentation](#) for more details.

Users can create additional network interfaces for the container if needed. Users can also give static IP addresses and routes to new interface.

By default, containers use the DNS settings from the host system. You can change the DNS policy and define separate nameservers and search domains. See the Docker [DNS services documentation](#) for more details.

Port Forwarding List

Choose the protocol and enter port numbers for both the container and node. Multiple port forwards can be defined. The node port number must be over 9000. Make sure no other containers or system services are using the same port number.

Host Path Volumes

Scale storage locations can be mounted inside the container. To mount Scale storage, define the path to the system storage and the container internal path for the system storage location to appear. You can also mount the storage as read-only to prevent the container from being used to change any stored data. For more details, see the [Kubernetes hostPath documentation](#).

Volumes

Users can create additional Persistent Volumes (PV's) for storage within the container. PV's consume space from the pool chosen for Application management. You will need to name each new dataset and define a path where that dataset appears inside the container.

To view created container datasets, go to **Storage** and expand the pool used for applications. Expand `/ix-applications/releases/<ContainerName>/volumes/ix-volumes/`.

Volume(s) Access

Users developing applications should be mindful that if an application uses Persistent Volume Claims (PVC), those datasets won't be mounted on the host, and therefore will not be accessible within a file browser. This is upstream `zfs-localpv` behavior which is being used for managing PVC(s)

If you want to consume or have file browser access to data that is present on the host, set up your custom application to use host path volumes.

Alternatively, you can use the network to copy directories and files to and from the pod using `k3s kubectl` commands.

To copy from a pod in a specific container: `k3s kubectl cp <file-spec-src> <file-spec-dest> -c <specific-container>`

To copy a local file to the remote pod: `k3s kubectl cp /tmp/foo <some-namespace>/<some-pod>:/tmp/bar`

To copy a remote pod file locally: `k3s kubectl cp <some-namespace>/<some-pod>:/tmp/foo /tmp/bar`

Deploying the Application

Saving an official or custom container adds a new entry to *Installed Applications*. The container enters a deploy status as it fetches the image from the remote repository and configures it. When deployment is complete, the container moves to an active status and can be used.



Accessing the Shell in an Active Container

To access the shell in an active container, first identify the namespace and pod for the container. In the Scale UI, go to **System Settings > Shell** to begin entering commands:

1. View container namespaces: `k3s kubectl get namespaces`.
2. View pods by namespace: `k3s kubectl get -n <NAMESPACE> pods`.
3. Access container shell: `k3s kubectl exec -n <NAMESPACE> --stdin --tty <POD> -- /bin/bash`.

Additional Container Commands

- View details about all containers: `k3s kubectl get pods,svc,daemonsets,deployments,statefulset,sc,pvc,ns,job --all-namespaces -o wide`.
- Get container status: `k3s kubectl describe -n <CONTAINER NAMESPACE> <POD-ID>`.

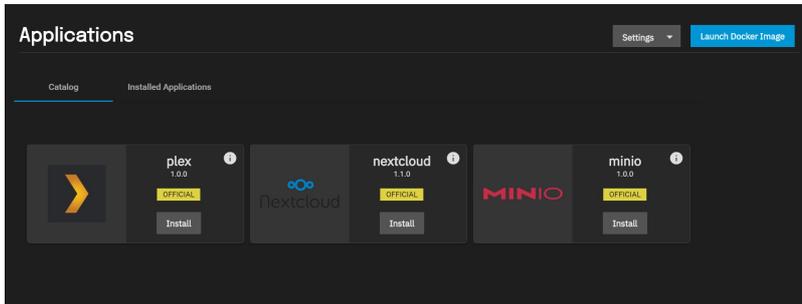
10.2 - Using SCALE Catalogs

- [Manage Catalogs](#)
 - [Adding Catalogs](#)
 - [Displaying Catalogs](#)

TrueNAS SCALE comes with a pre-built official catalog of iXsystems-approved Docker apps that includes Plex, [MinIO](#), Nextcloud, [Chia](#), and IPFS. Users can also configure custom apps catalogs, although iXsystems will not directly support any non-official apps in a custom catalog.

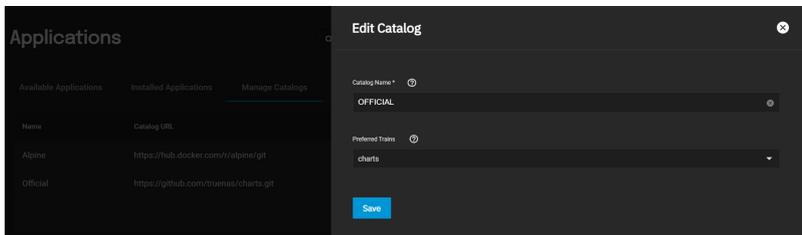
Manage Catalogs

To manage and add catalogs, click the *Manage Catalogs* tab in the **Applications** screen.



Users can edit, refresh, delete, and view the summary of a catalog by clicking the  button next to the intended catalog.

Edit: The edit option allows users to respecify the name TrueNAS will use to look up the catalog, as well as the trains from which UI should retrieve available applications for the catalog.

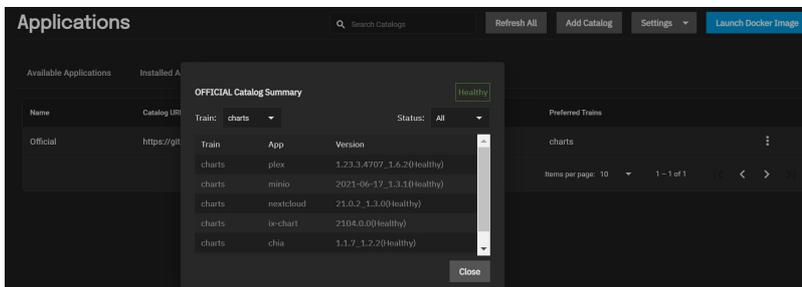


Refresh: The Refresh option re-pulls the catalog from its repository and applies any updates.

Delete: The Delete option allows users to remove a catalog from the system. Users cannot delete the default Official catalog.

Summary: The Summary option lists all of the apps in the catalog and sorts them Train, App, and Version.

Users may filter the list by Train type (All, charts, or test), and by Status (All, Healthy, or Unhealthy).



Adding Catalogs

To add a catalog, click the *Add Catalog* button in the *Manage Catalogs* tab and fill out the form. As an example, fill out the form using the data below to add the Truecharts catalog to SCALE.

Field	Description	Truecharts
Catalog Name	The name the TrueNAS will use to look up the catalog.	truecharts
Force Create	Set to add the catalog to the system even if some trains are unhealthy.	Unchecked
Repository	The valid git repository URL.	https://github.com/truecharts/catalog
Preferred Trains	The trains TrueNAS will use to retrieve available applications for the catalog.	stable (and optionally: incubator)
Branch	The git repository branch TrueNAS will use for the catalog.	main

Displaying Catalogs

Users can select which catalogs they want to view in the *Available Applications* tab by clicking the *Catalogs* drop-down menu and checking the catalogs they want to see.

Applications

Search Available Applications

Refresh All

Catalogs

Settings

Launch Docker Image

- Select All
- Official
- Truecharts

Available Applications

Installed Applications

Manage Catalogs

Manage Docker Images

 Official Charts Install	plex 1.23.3.4707_1.6.2 Install	 Official Charts Install	minio 2021-06-17_1.3.1 Install	 Official Charts Install	nextcloud 21.0.2_1.3.0 Install
 Official Charts Install	chia 1.1.7_1.2.2 Install	 Official Charts Install	ipfs v0.9.0_1.1.1 Install	 Truecharts Stable Install	zwavejs2mqtt auto_0.2.11 Install
 Truecharts Stable Install	syncthing auto_6.2.10 Install	 Truecharts Stable Install	jackett auto_6.2.11 Install	 Truecharts Stable Install	vaultwarden auto_3.3.11 Install

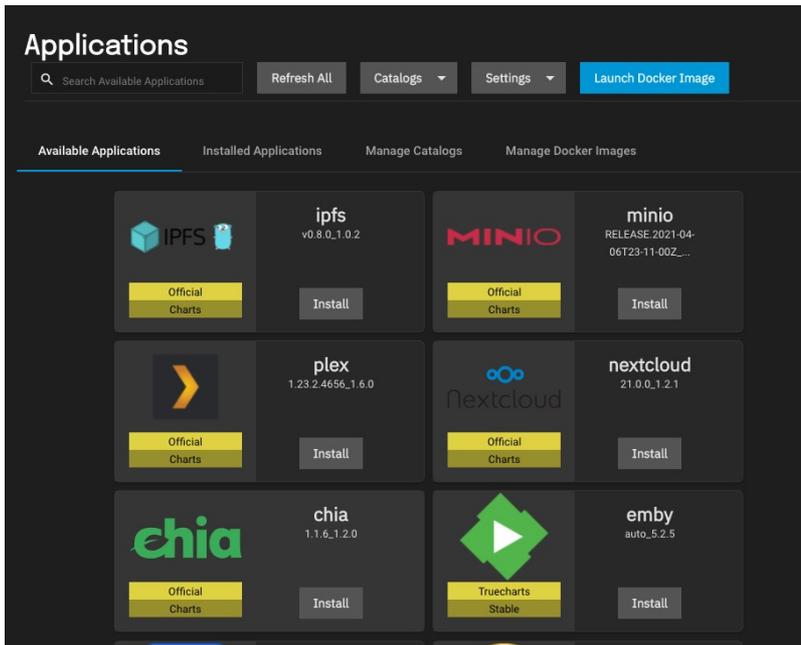
10.3 - Chia App

- [install the Chia App](#)

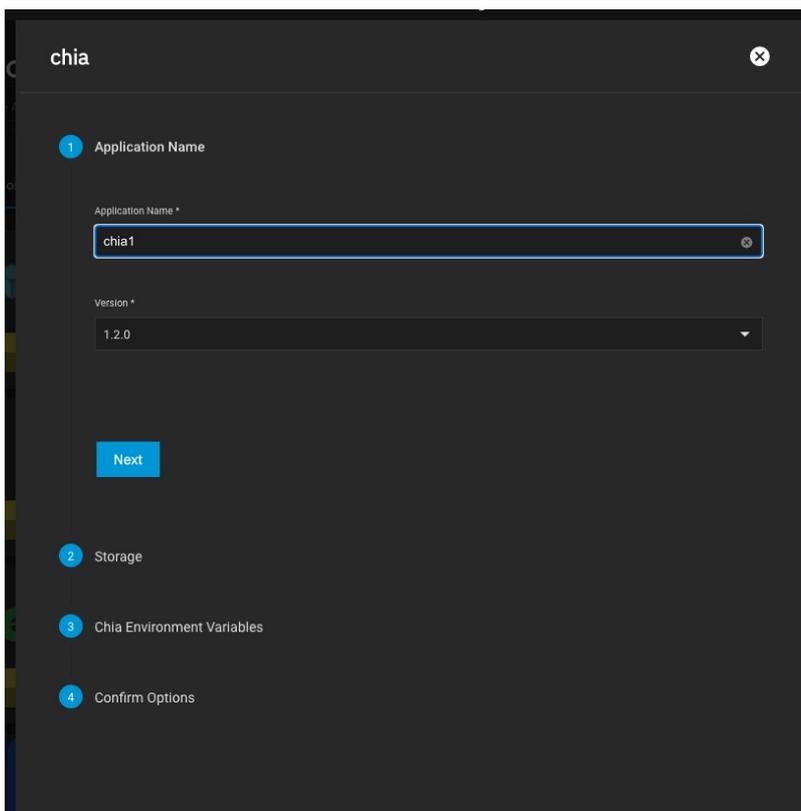
SCALE includes Chia in its Official Apps catalog. Chia Blockchain is a new cryptocurrency that uses Proof of Space and Time. Instead of using expensive hardware that consumes exorbitant amounts of electricity to mine cryptos, it leverages existing empty hard disk space on your computer(s) to farm cryptos with minimal resources, such as electricity.

install the Chia App

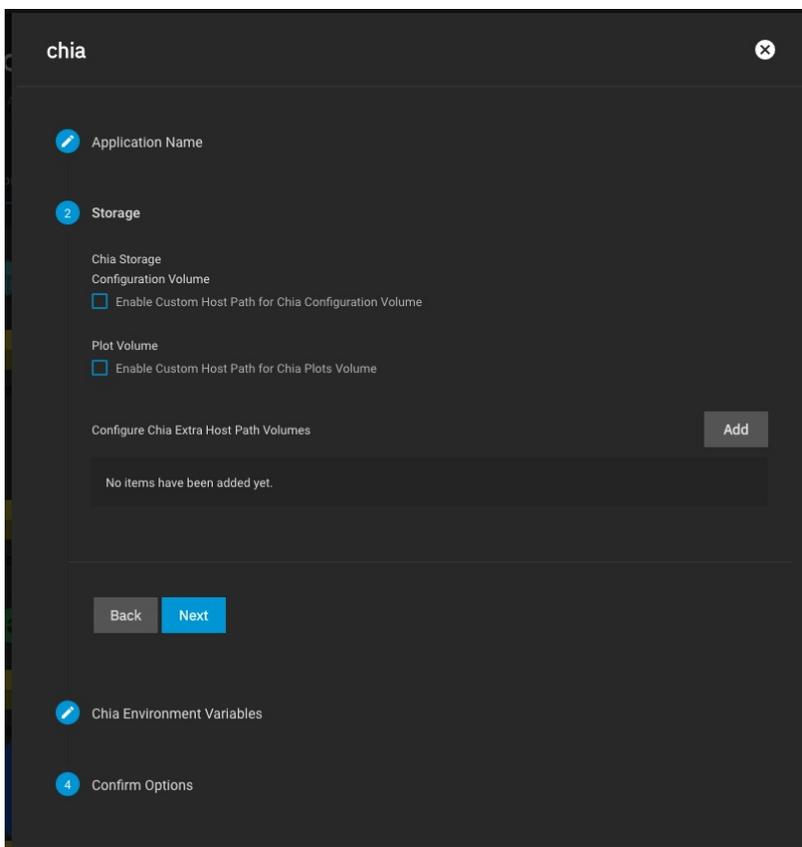
Click on the Chia app's **Install** button in the *Available Applications* list.



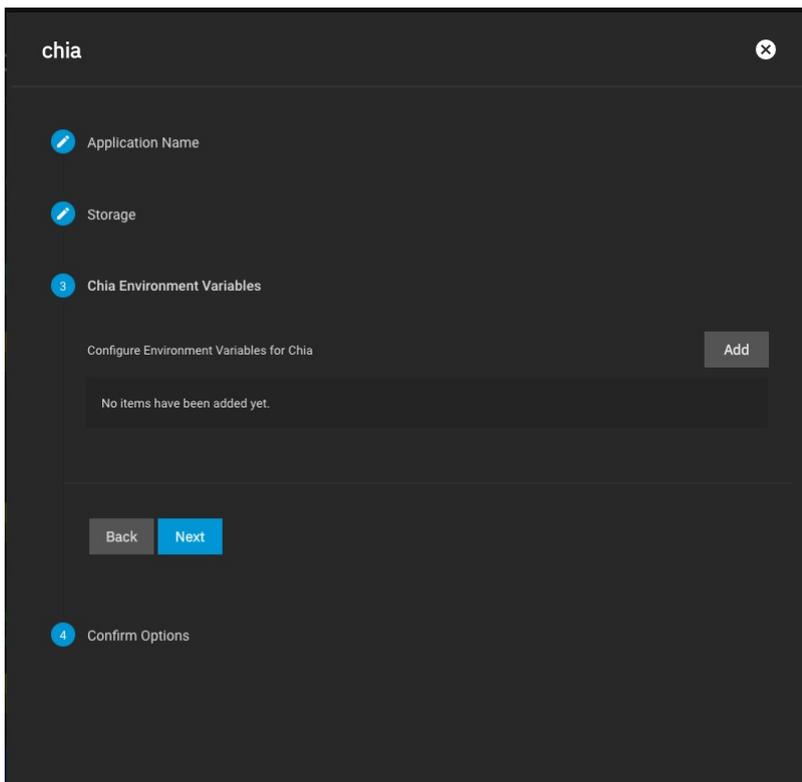
Name your App and click *Next*. In this example, the name is *chia1*.



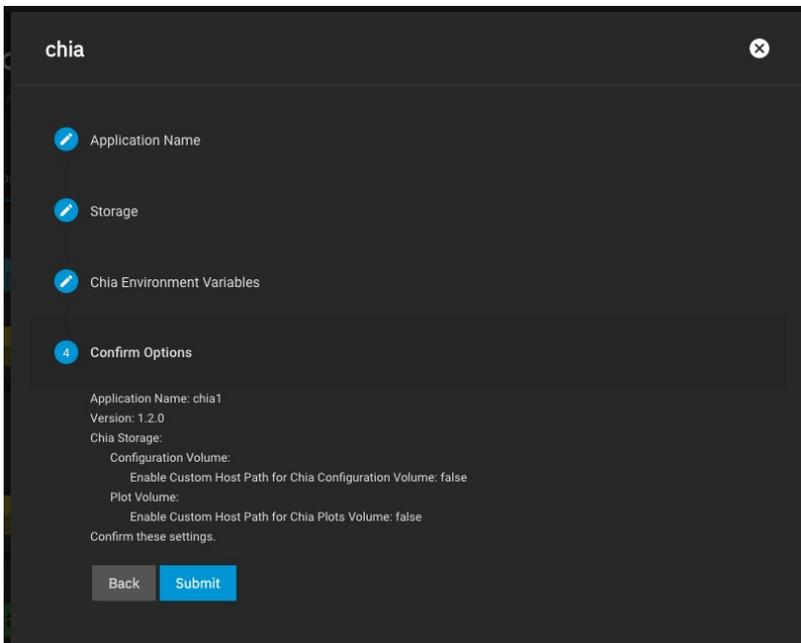
Leave *Enable Custom Host Path for Chia Configuration Volume* and *Enable Custom Host Path for Chia Plots Volume* unchecked and click *Next*.



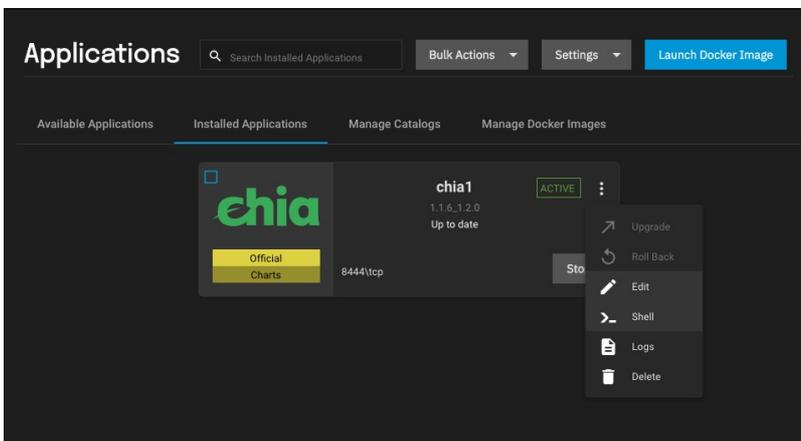
Click *Next* in the Chia Environment Variables screen. You will add one later.



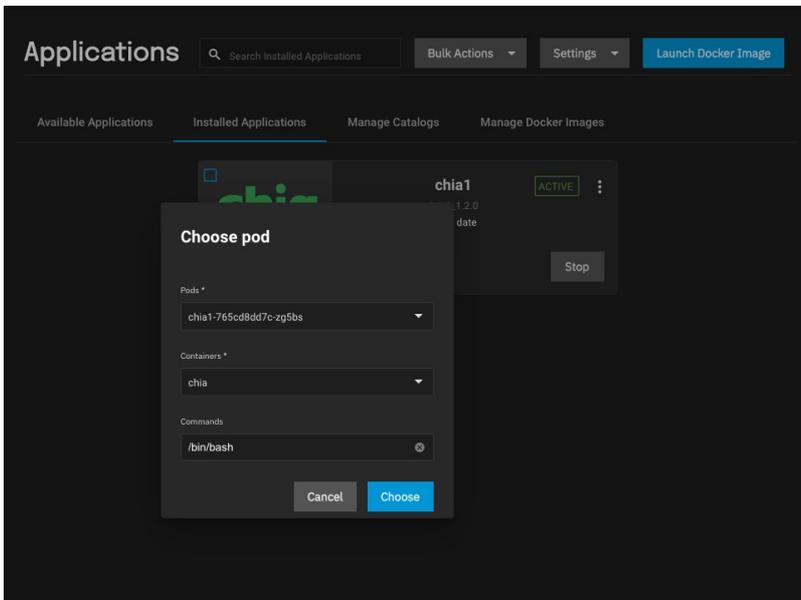
Confirm the options and click *Submit*.



Continue through the wizard and create the new application. After a minute or two the new Chia container will start and show "ACTIVE" status. Click the three-dot menu on the top-right and launch the Shell.



Leave the defaults for the pod (there is only one) and use the selected /bin/bash shell.



The first time Chia launches, it will automatically create a new private key set (for plotting purposes) and wallet. However, the private key set will not be preserved across container restarts. To make sure your keys and wallet persist, you need to save the Mnemonic Seed that was created and make sure it gets used at each container initialization. To do this, start by displaying the current key information by running the following shell command:

```
/chia-blockchain/venv/bin/chia keys show --show-mnemonic-seed
```

```
root@scaley:/chia-blockchain# /chia-blockchain/venv/bin/chia keys show --show-mnemonic-seed
Showing all public and private keys

Fingerprint: 3599713715
Master public key (m): ae033bd2c021c0591925bc1f0791db011f506c32965e4395e56c2a7e332ace734e1c6f9bdd6ccb75832088b7d1e961a
Farmer public key (m/12381/8444/0/0): azd6449467ce52092c14fe2f9f8e944448c15f8d68f7f672ac1baf8c7ce0866db20db1ac4a14b7f3fed1b9a32cd31
Pool public key (m/12381/8444/1/0): 8f8444b2372716d31e417f0ba9d64d7d8be291b5ca4379f4d9ed87ae53ab093d9e6737f3bdc8cb2d3e00c79d2e32b6
First wallet address: xch1hga4d08sgwyseu8mju8j8uhmap87ymd2n3dmhpr5zduu2k5j3pqwf4y6
Master private key (m): 12c58ccdda86836502244e4f355a18fa3fbae52bcecb2e98b99f2bb1e8ee04
First wallet secret key (m/12381/8444/2/0): 928b8a1f2bbd2b04fccf44b516176094a2a8c6b8f8c1dd26b57aad86ee18d68ec77fc1ddfe490c5e410833c6d4083f
Mnemonic seed (24 secret words):
spare seat media winter what remain river roast hard bean legend script luggage snack token real supreme scan uncle beef play a
ware senior garment
root@scaley:/chia-blockchain#
```

We suggest you make a backup copy of the information provided here for your reference in case you lose the keyfile. To make sure the same key is used for this container going forward, you will be saving the mnemonic-seed phrase to one of your host volumes on TrueNAS.

Copy and paste the 24 secret words of the mnemonic seed into a new shell command:

```
echo "my unique 24 secret words here" > /plots/keyfile
```

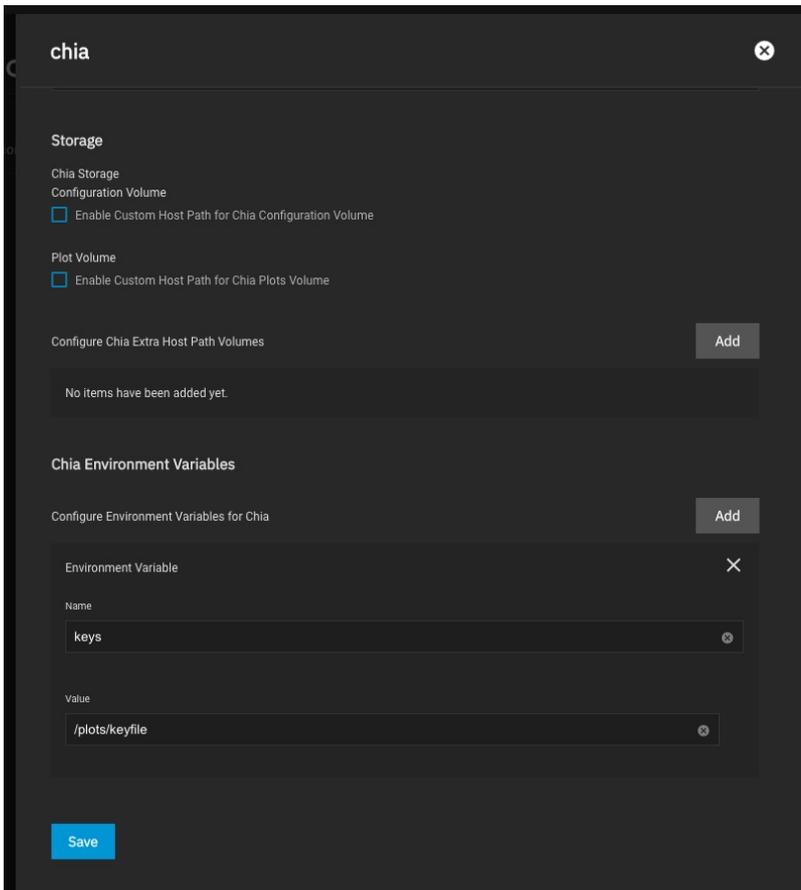
```
root@scaley:/chia-blockchain# /chia-blockchain/venv/bin/chia keys show --show-mnemonic-seed
Showing all public and private keys

Fingerprint: 3599713715
Master public key (m): ae033bd2c021c0591925bc1f0791db011f506c32965e4395e56c2a7e332ace734e1c6f9bdd6ccb75832088b7d1e961a
Farmer public key (m/12381/8444/0/0): azd6449467ce52092c14fe2f9f8e944448c15f8d68f7f672ac1baf8c7ce0866db20db1ac4a14b7f3fed1b9a32cd31
Pool public key (m/12381/8444/1/0): 8f8444b2372716d31e417f0ba9d64d7d8be291b5ca4379f4d9ed87ae53ab093d9e6737f3bdc8cb2d3e00c79d2e32b6
First wallet address: xch1hga4d08sgwyseu8mju8j8uhmap87ymd2n3dmhpr5zduu2k5j3pqwf4y6
Master private key (m): 12c58ccdda86836502244e4f355a18fa3fbae52bcecb2e98b99f2bb1e8ee04
First wallet secret key (m/12381/8444/2/0): 928b8a1f2bbd2b04fccf44b516176094a2a8c6b8f8c1dd26b57aad86ee18d68ec77fc1ddfe490c5e410833c6d4083f
Mnemonic seed (24 secret words):
spare seat media winter what remain river roast hard bean legend script luggage snack token real supreme scan uncle beef play a
ware senior garment
root@scaley:/chia-blockchain# echo "spare seat media winter what remain river roast hard bean legend script luggage snack token
real supreme scan uncle beef play aware senior garment" > /plots/keyfile
root@scaley:/chia-blockchain#
```

Now exit the shell and go back to the *Installed Apps* page. Click *Edit* on your Chia container.

Scroll down until you find the Container Environment Variables section and add a new variable as shown below:

- Environment Variable Name: keys
- Environment Variable Value: /plots/keyfile



If you entered the command correctly, you should see some output that looks like the screenshot.

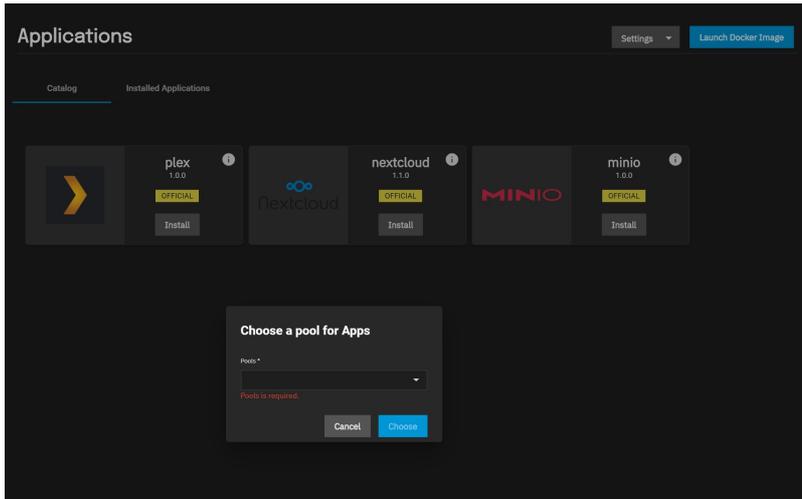
Save the change, and the chia container should restart automatically. To confirm your changes have persisted you can log into the containers shell again and run the same `/chia-blockchain/venv/bin/chia keys show --show-mnemonic-seed` command to show your keys. If the keys are identical to what you previously recorded, then you are done! This Chia container will persist across reboots, upgrades, and re-deployments.

At this point, you are ready to begin farming Chia. This is a CLI process and beyond the scope of this quick how-to, but we recommend you start by reading up on their [CLI reference materials](#), [Quick Start guide](#) and other [documentation](#).

10.4 - Deploying TrueCommand on TrueNAS SCALE

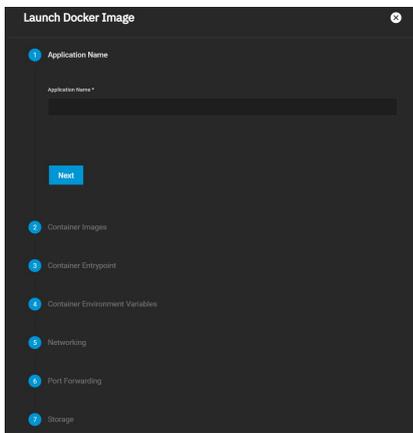
A [TrueCommand](#) Docker image can be deployed on TrueNAS SCALE. The SCALE installation needs internet access and a storage pool with at least 20 GiB of space available for the image.

After logging in to SCALE, go to the **Apps** page. The first time the Apps page opens, select an existing pool to store any installed Applications. It is recommended to choose a pool that has enough space available for all the different applications you intend to install and maintain. Select the desired pool from the drop down and click **Choose**.



This can be changed later by opening the *Settings* drop down and clicking *Choose Pool*.

After selecting a pool, click *Launch Docker Image*.



This opens a wizard to install a custom application from a Docker image repository.

Use lowercase alphanumeric characters for the application name and enter `ixsystems/truecommand` for the *Image Repository*. Enter `latest` for the *Image Tag*.

Install Application



1 Image and Policies

Application Name *

truecommand01

Image Repository *

ixsystems/truecommand

Image Tag

latest

Image Pull Policy

Only pull image if not present on host.

Update Strategy

Create new pods and then kill old ones.

Restart Policy

Always restart containers in a pod if they exit.

Next

2 Container Settings

3 Networking

4 Port Forwarding List

5 Host Path Volumes

6 Volumes

While not required, you can customize the *Container Settings* according to your use case or environment.

2 Container Settings

Container CMD ⓘ

Container Args ⓘ

Add containerEnvironmentVariables Add

Environment Variable Name ×

TZ ×

Environment Variable Value ×

America/NewYork ×

Environment Variable Name ×

WEBPASSWORD ×

Environment Variable Value ×

s3cur3p4\$\$w0rd ×

Back Next

As before, modifying the *Networking Settings* is not required, but you might need to customize according to your environment.



Image and Policies

Container Settings

3 Networking

Host Network

Add External Interface



Host Interface

IP Address Management

Static IP

 / 24

Add

Static Route: Destination

 / 24

Static Route: Gateway

Add

DNS Policy

DNS Configuration

Nameservers

Searches

Back

Next

To access the TrueCommand interface after installing the application, add one or two entries to the *Port Forwarding List*. For HTTP access to TrueCommand, enter *80* for the *Container Port* and *9004* for the *Node Port*.

To include HTTPS access, click the + button to add another port forwarding dialog. In the second form, enter *443* for the *Container Port* and *9005* for the *Node Port*.

Install Application



1 Image and Policies

2 Container Settings

3 Networking

4 Port Forwarding List

Container Port

80

Node Port

9004

Protocol

TCP Protocol

Container Port

443

Node Port

9005

Protocol

TCP Protocol

Back

Next

5 Host Path Volumes

6 Volumes

7 Confirm Options

You can link additional SCALE storage to this application by adding *Host Path Volumes* entries.

Install Application



1 Image and Policies

2 Container Settings

3 Networking

4 Port Forwarding List

5 Host Path Volumes

Host Path



+

▶ /mnt

Mount Path ⓘ

Read Only

Back

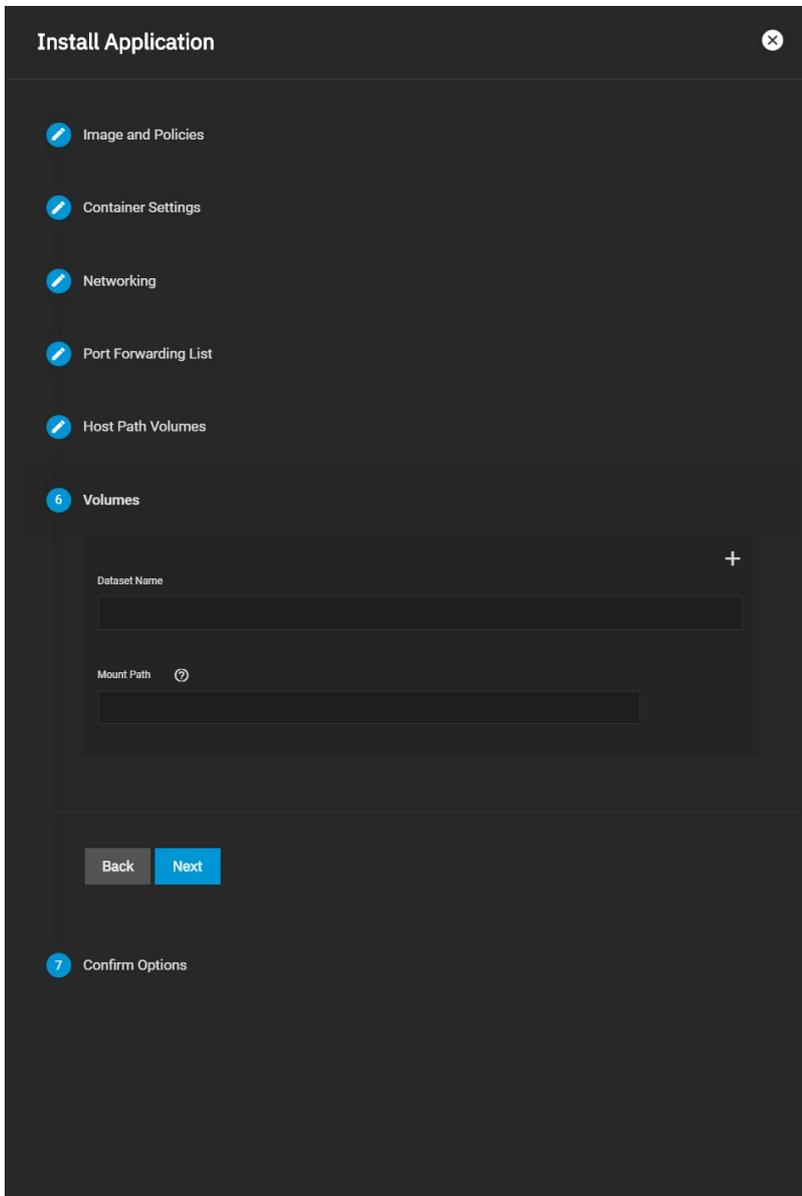
Next

6 Volumes

7 Confirm Options

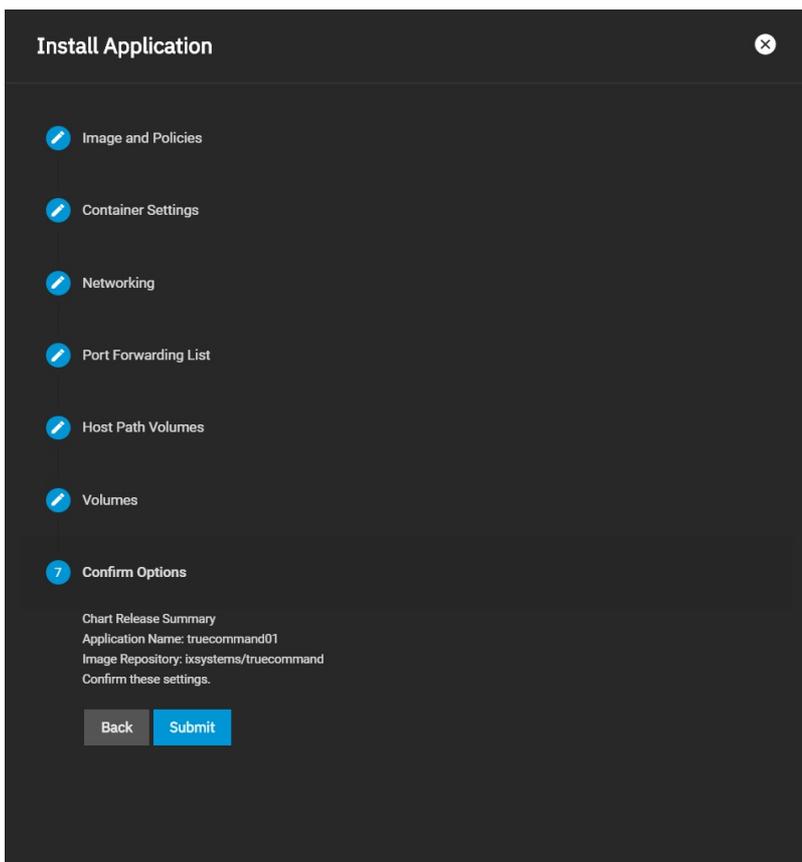
This is not typically required for TrueCommand.

Volumes is for creating additional storage datasets inside the previously selected Applications Pool.

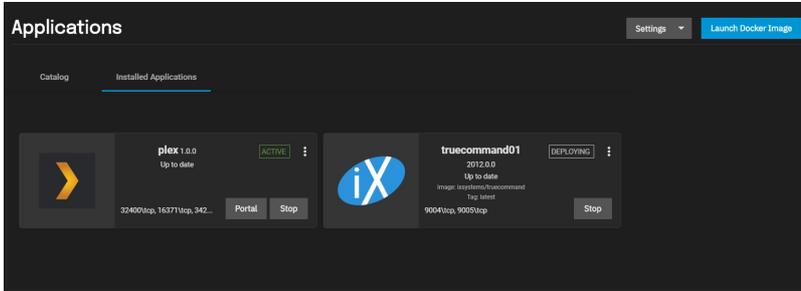


This is not typically required for TrueCommand.

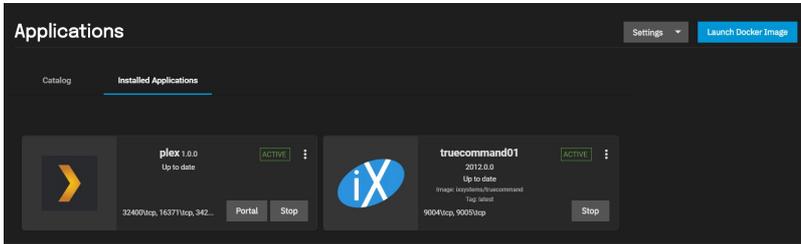
Review the settings and verify the *Image Repository* is correct.



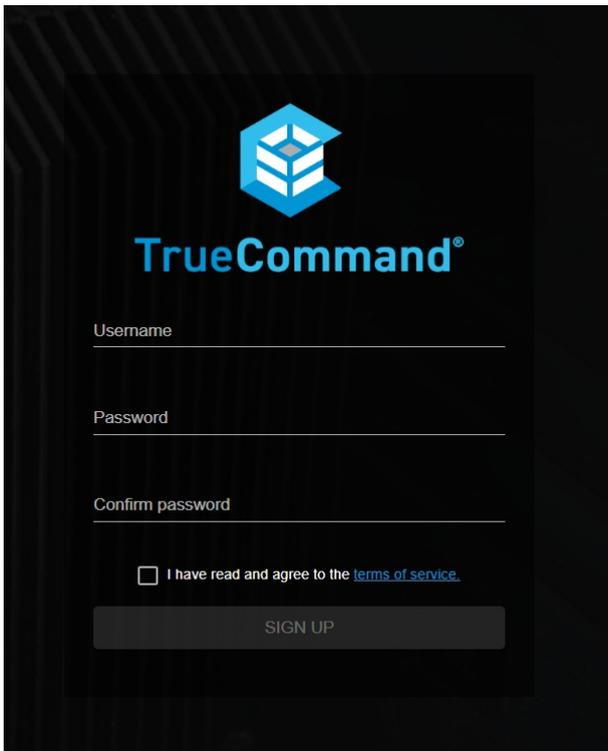
Click **Submit** to download and begin deploying TrueCommand. A new card is added to **Installed Applications** and shows the container deployment status.



When deployment is complete, the card updates to show TrueCommand is active.



To access TrueCommand, open a new browser tab and enter the address of your TrueNAS SCALE with :9004 or :9005 appended. Example: <https://www.truenasscale01.ixsystems.com:9005> The TrueCommand login screen appears and asks to create the new [Administrator account](#).



If the login screen fails to appear, double-check your system networking settings, open networking ports, and if the 9004 or 9005 *Node Port* values are already in use by another application.

For more details about TrueCommand, see the [TrueCommand section](#).

10.5 - Using Docker on TrueNAS SCALE

SCALE includes the ability to run Docker containers using Kubernetes.

What is Docker?

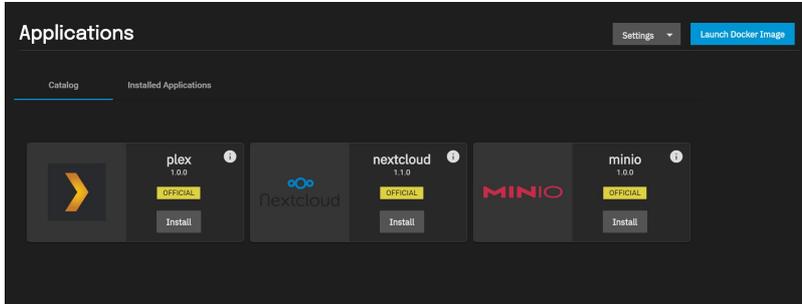
Docker is an open platform for developing, shipping, and running applications. Docker enables the separation of applications from infrastructure through OS-level virtualization to deliver software in containers.

What is Kubernetes?

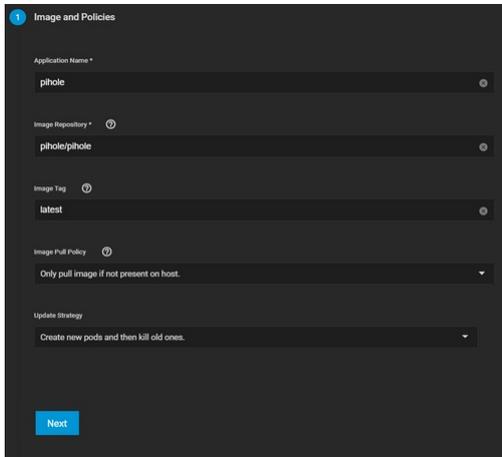
Kubernetes is a portable, extensible, open-source container-orchestration system for automating computer application deployment, scaling, and management with declarative configuration and automation.

Always read through the Docker Hub page for the container you are considering installing so that you know all of the settings that need to be configured. To Set up a Docker Image, first determine if you wish the container to use its own dataset. Create a dataset for before hand if desired for host volume paths.

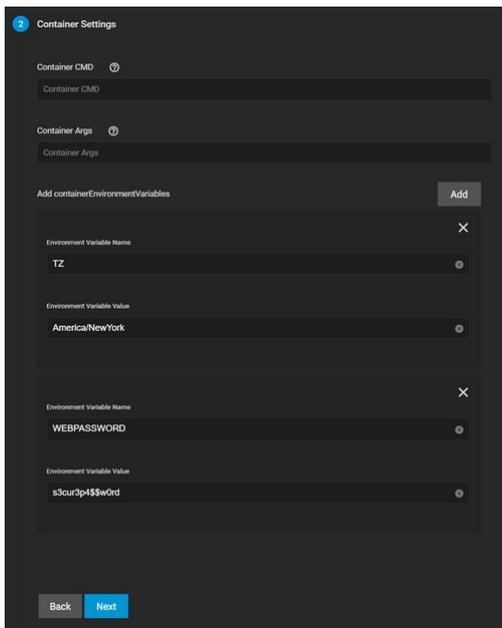
When you are ready to create a container, open the **APPS** page and click **Launch Docker Image**.



Fill in the *Application Name* and *Image Repository* for the docker container you'd like to set up. This example shows installation of the [PiHole project](#).



Click **Next** to move to the *Container Settings* section. In this example, PiHole needs the timezone and password set. Always refer to the dockerhub page for information on what the docker container requires.



Clicking **Next** will open the Networking section. If the container needs special networking configuration it should be set here. Ports are configured in the next section. Once completed, click Next to move forward in the configuration process.

The PiHole Docker Hub page lists a set of ports that will need to be set. These values may need to be adjusted depending on the configuration of your system. TrueNAS SCALE requires all Node Ports to be above 9000.

Click **Next** when all the ports are configured.

The Host Path volume will be set to a dataset and directory paths which were created before attempting to deploy the container. PiHole uses volumes store your data between container upgrades. You will need to create these directories in a dataset on SCALE prior to installing this container.

Additional Volumes can be added to the container if needed. When all the settings have been entered, verify the Application and Container Name and click **Submit**.



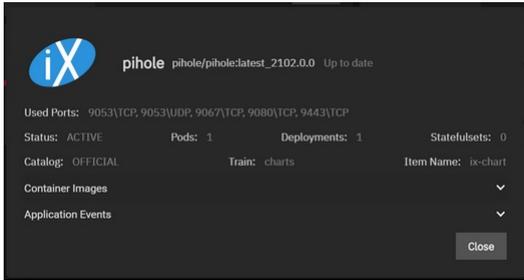
TrueNAS SCALE will deploy the container.



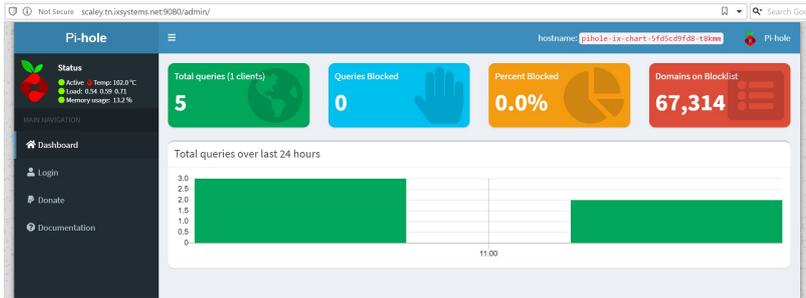
When the deployment is completed the container will become active. If the container does not autostart, click the **START** button.



Clicking on the App card will reveal details.



With PiHole as our example we navigate to the IP of our TrueNAS system with the port and directory address :9080/admin/.



10.6 - MinIO Clusters

- [First Steps](#)
- [Configuring MinIO](#)

On TrueNAS SCALE 20.12-ALPHA and later, users can create a MinIO S3 distributed instance to scale out and handle individual node failures. A node refers to a single TrueNAS storage system in a cluster.

In the images below, we used four TrueNAS systems to create a distributed cluster. For more information on MinIO distributed setups, refer to the [MinIO documentation](#).

First Steps

Before you configure MinIO, you must create a dataset and shared directory for the persistent MinIO data. Go to **Storage > Pools** and select the pool you want to place the dataset in. You can use an existing pool or create a new one.

After creating the dataset, go to **System > Shell** and create the directory MinIO will store information the application uses. MinIO uses **/data** but allows users to replace this with the directory of their choice. Change to the `/pool/dataset` directory and then use the `mkdir /mnt/data` command to create the **/data** directory.

For a distributed configuration, repeat this on all system nodes in advance.

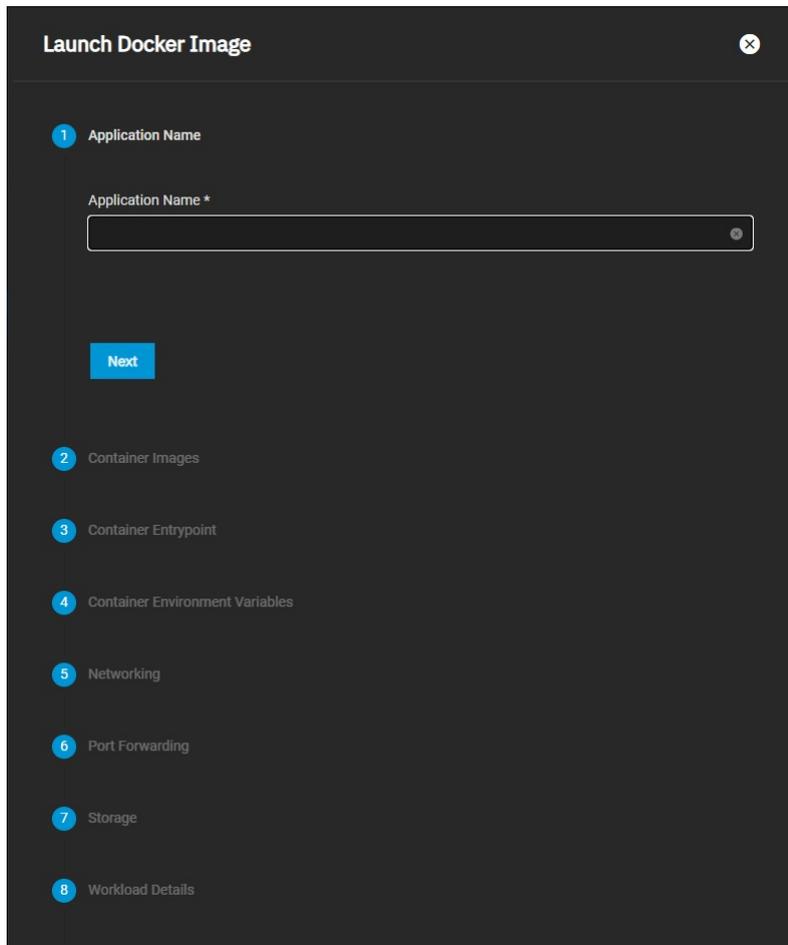
Note the system (node) IP addresses or hostnames and have them ready for configuration. Also, have your S3 username and password ready for later.

Configuring MinIO

You can configure the MinIO application using either the **Launch Docker Image** button or the **Install** button on the MinIO application card on the **Available Applications** tab.

Setup Using Launch Docker Image

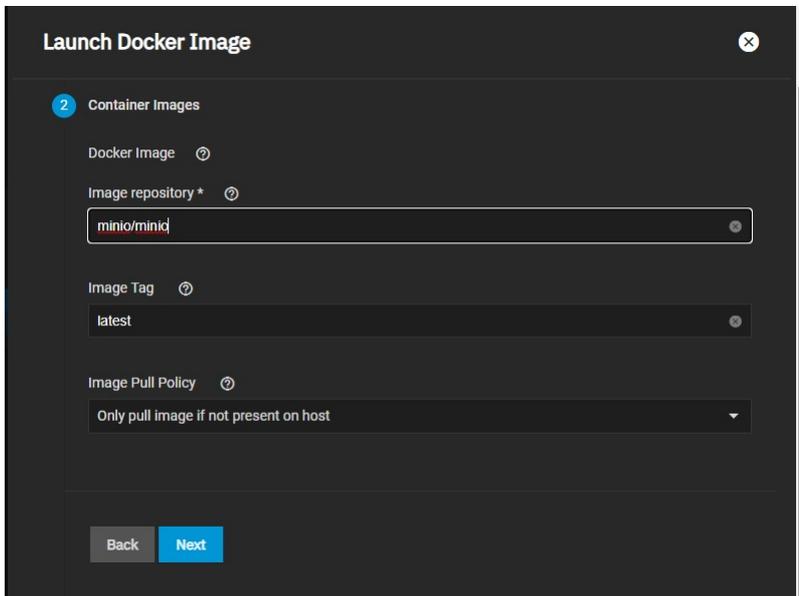
On your first node, go to **Apps** and click **Launch Docker Image**.



First, enter a name in **Application Name** (for example, `minio` for a normal configuration or `minio-distributed` for a distributed MinIO configuration). A MinIO in distributed mode allows you to pool multiple drives (even if they are different machines) into a single object storage server for better data protection in the event of single or multiple node failures because MinIO distributes the drives across several nodes. For more information, see the [Distributed MinIO Quickstart Guide (<https://docs.min.io/docs/distributed-minio-quickstart-guide>)].

Click **Next** to continue after completing each section of the configuration form.

Enter `minio/minio` as the image name under **Image Repository**. Click **Next**.

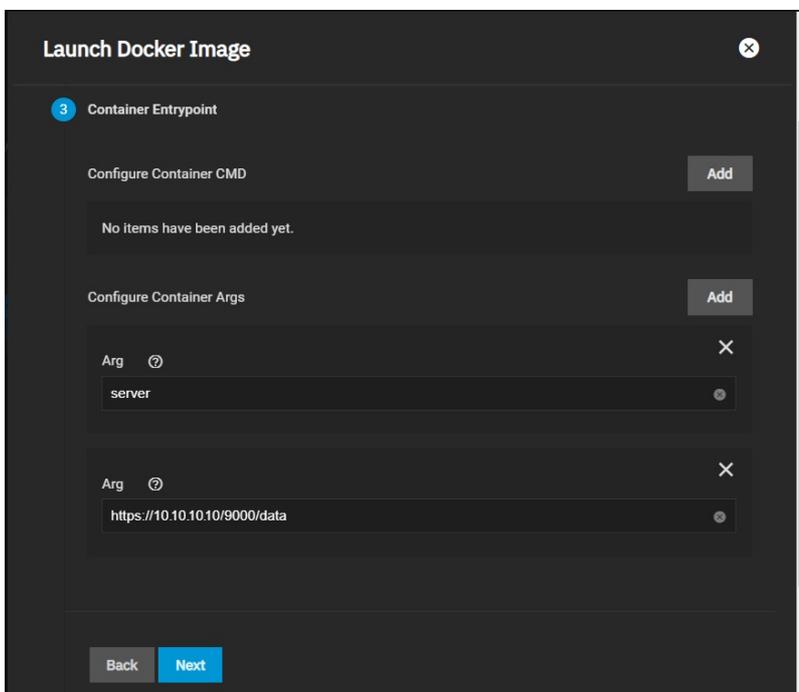


Configure the **Container Entrypoint** arguments. Click the **Add** button to the right of **Configure Container Args** twice to add two **Arg** fields. In the first **Arg** field type **server**. In the second **Arg** field, type the valid IP or hostname of each TrueNAS system on the network, the MinIO port number, and the directory you created for MinIO. Use this format: **http://0.0.0.0/9000/data**.

For a distributed cluster, add the valid TrueNAS system (node) IP addresses/hostnames. The order is important, so use the same order across all the nodes.

MinIO containers use server port 9000. The MinIO Console communicates using port 9001.

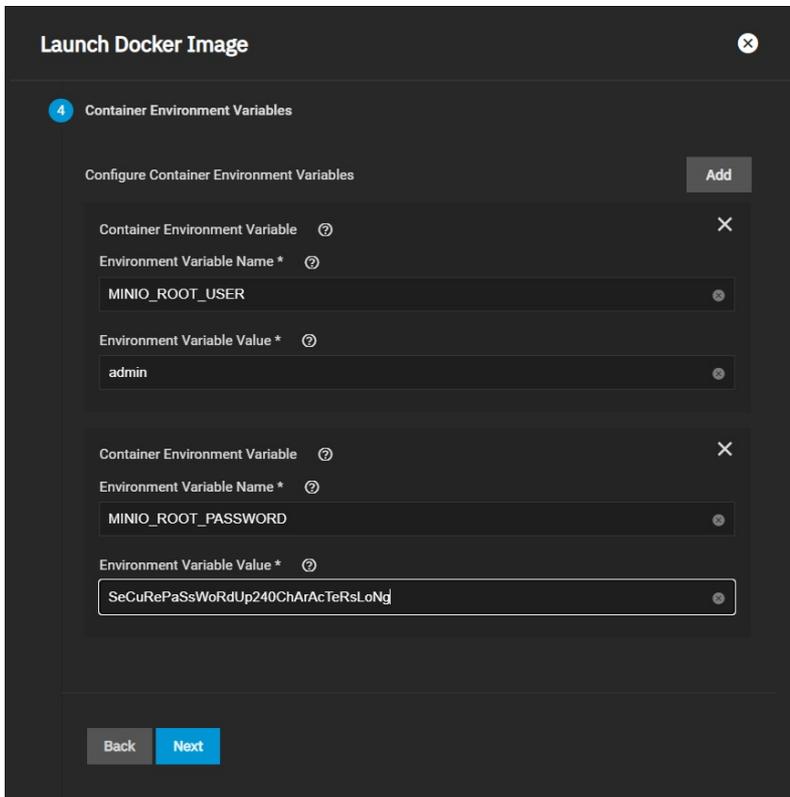
Use the /data path which is set up in the next steps.



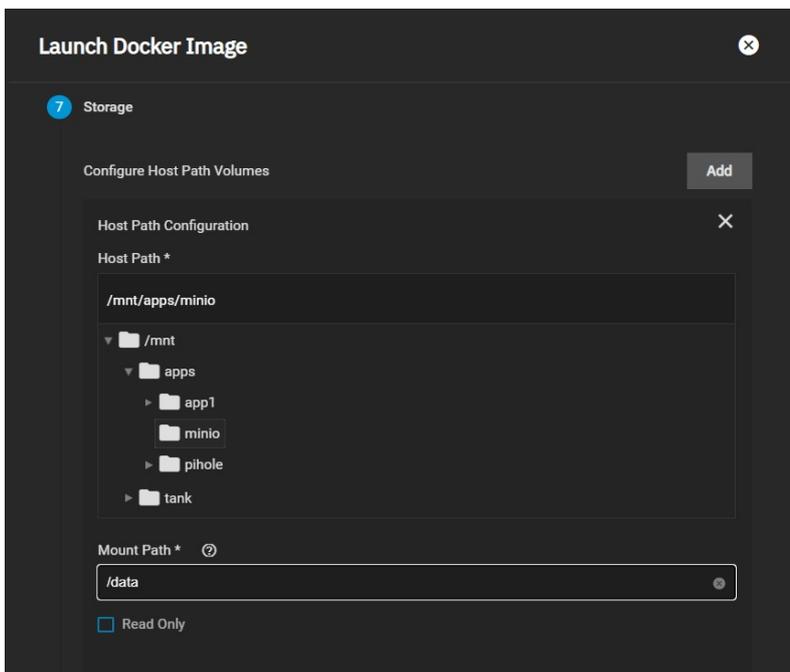
Next, create the **Container Environment Variables** and define the **MINIO_ROOT_USER** and **MINIO_ROOT_PASSWORD** arguments and their values. For the **ROOT_USER** value, use a name up to 20 characters. For the **ROOT_PASSWORD**, use a string of 8 to 40 randomized characters. MinIO recommends using a long password string of unique random characters. Refer to [MinIO User Management](#) for more information.

Keep all passwords and credentials secured and backed up.

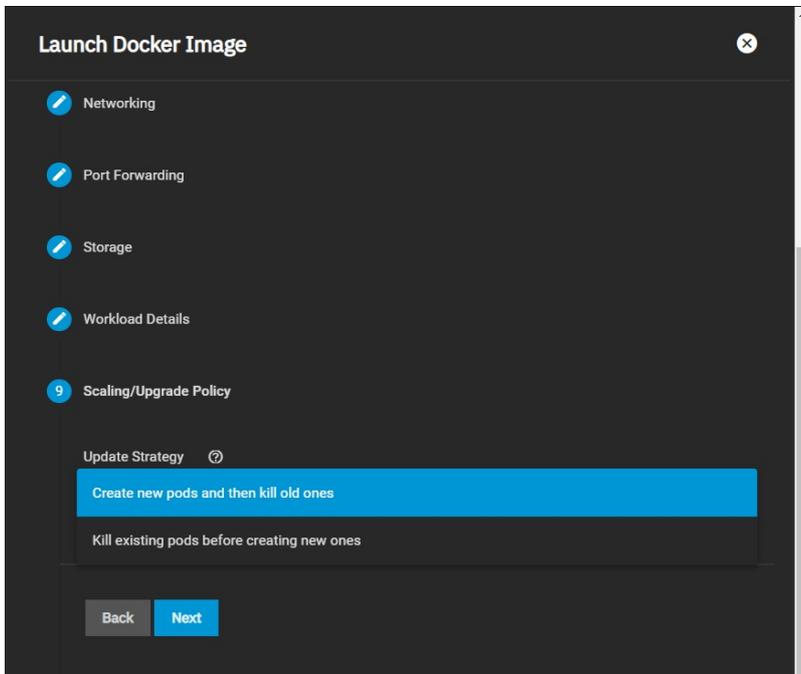
For a distributed cluster, ensure the values are identical between nodes and fill the **Environment Variable Value** with proper random credentials.



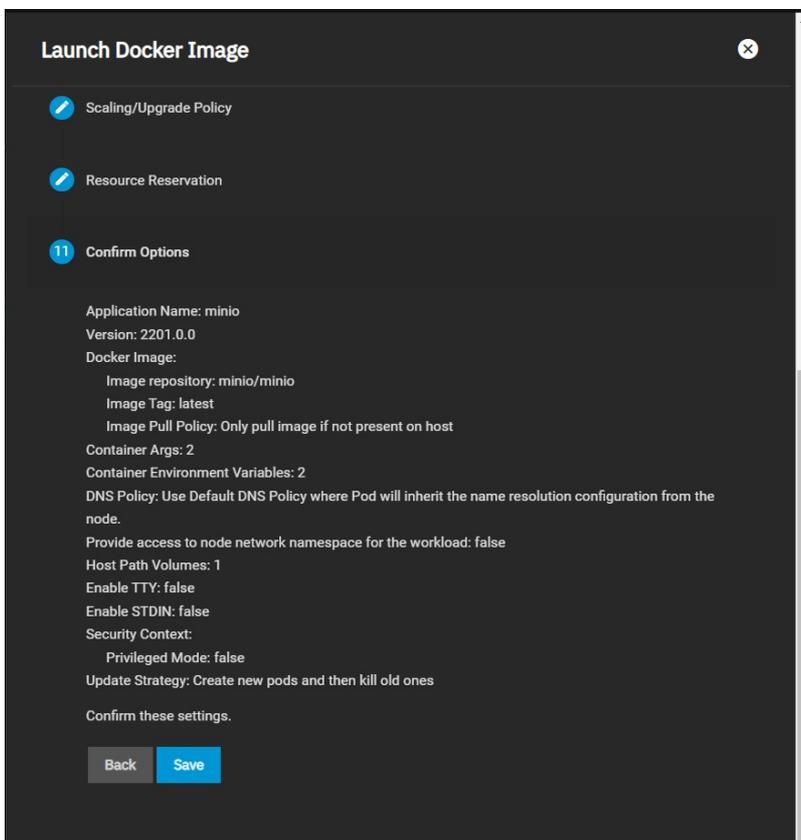
Click **Next** until the **Storage** section displays. Select the dataset you created for the MinIO container for the **Host Path** and enter the **/data** directory under **Mount Path**, then click **Next**.



Click **Next** until you reach the **Scaling/Upgrade Policy** screen. Select the **Update Strategy** option you want to deploy. Use **Kill existing pods before creating new ones** to recreate the container or **Create new pods and then kill old ones** if you want rolling upgrades. Click **Next**.



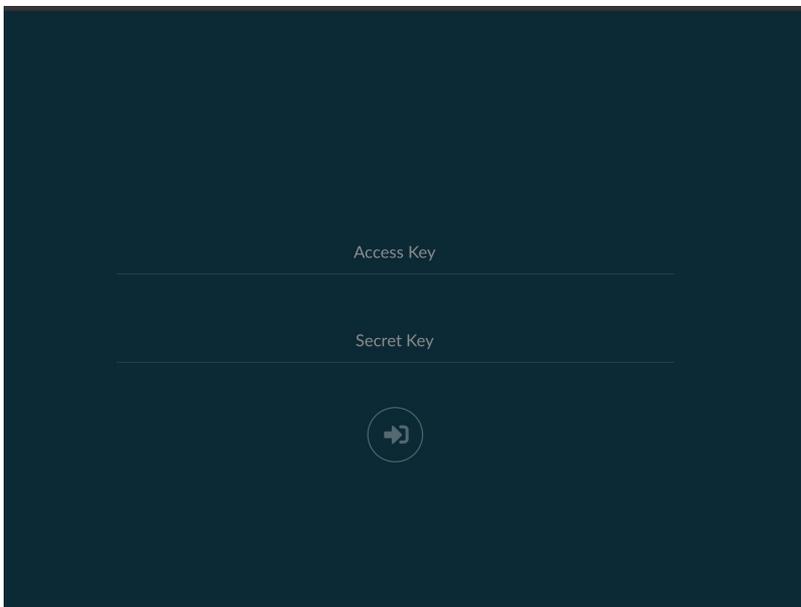
Confirm your options, then click **Save** to complete the first node.



Now that the first node is complete, you can configure any remaining nodes (including datasets and directories).

Accessing the Minio Setup

Once you're done creating datasets, you can navigate to the TrueNAS address at port **:9000** to see the MinIO UI. If you created a distributed setup, you can see all your TrueNAS addresses. Log in with the **ROOT_USER** and **ROOT_PASSWORD** keys you created as Container Environment Variables.



Setting Up Using MinIO Install

Go to **Apps** and select the **Available Applications** tab to display the MinIO application card. Click **Install** on the MinIO card to open the MinIO configuration wizard.

Application Name

Enter a name for the MinIO cluster. Click **Next**. Type the name in all lowercase.

A dark-themed configuration window titled "minio" with a close button (X) in the top right. The first step is "1 Application Name". It contains a text input field for "Application Name *" and a dropdown menu for "Version *" with "1.5.2" selected. A blue "Next" button is at the bottom left.

Workload Configuration

Select an update strategy. Use **Kill existing pods before creating new ones** to recreate the container or **Create new pods and then kill old ones** if you want rolling upgrades. We recommend **Kill existing pods before creating new ones**. Click **Next**.

A dark-themed configuration window titled "minio" with a close button (X) in the top right. The second step is "2 Workload Configuration". It shows "Minio update strategy" with two radio button options: "Create new pods and then kill old ones" (which is selected and highlighted in blue) and "Kill existing pods before creating new ones". At the bottom are "Back" and "Next" buttons.

MinIO Configuration

If you want to run your MinIO instance to connect to a distributed MinIO cluster, set **Enable Distributed Mode** and input your Distributed Minio Instance URI. See the [Distributed MinIO Quickstart Guide](https://docs.min.io/docs/distributed-minio-quickstart-guide) for more information.

minio

3 Minio Configuration

Enable Distributed Mode ⓘ

Configure Minio Extra Arguments Add

No items have been added yet.

Root User * ⓘ

Root Password * ⓘ

Configure Minio image environment Add

No items have been added yet.

Minio Service Configuration ⓘ

Node Port to use for Minio API *

9000

Node Port to use for Minio UI Access *

9002

Minio Domain Name ⓘ

Click the **Add** button to the right of **Configure MinIO Extra Arguments** twice to display two **Arg** fields. In the first **Arg** field type **server**. In the second **Arg** field type the valid IP or hostname of each TrueNAS systems on the network, the MinIO port number, and the directory you created for MinIO. Use this format, **http://0.0.0.0/9000/data**.

Add the other valid TrueNAS system IP addresses/hostnames of your various nodes. The order is important, so use the same order across all the nodes. MinIO containers use server port 9000. The MinIO UI communicates using port 9002.

minio

3 Minio Configuration

Enable Distributed Mode ⓘ

Configure Minio Extra Arguments Add

Argument ×

server

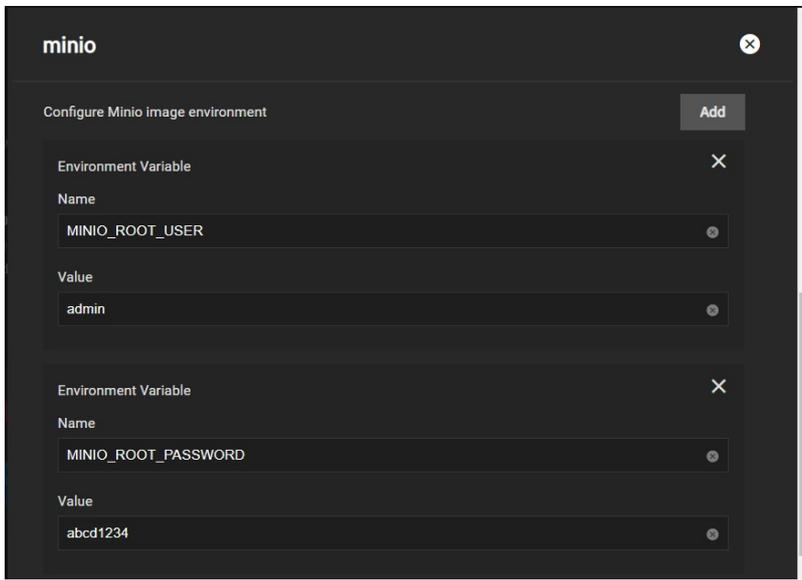
Argument ×

http://10.10.10.10/9000/data

Enter the S3 root user in **Root User** and the S3 password in the **Root Password** fields.

Click the **Add** button to the right of **Container Environment Variables** and enter the **MINIO_ROOT_USER** and **MINIO_ROOT_PASSWORD** arguments and values. For the **ROOT_USER** value, use a name up to 20 characters. For the **ROOT_PASSWORD**, use 8 to 40 randomized characters. MinIO recommends using a long password string of unique random characters. Refer to [MinIO User Management](#) for more information.

Keep all passwords and credentials secured and backed up.

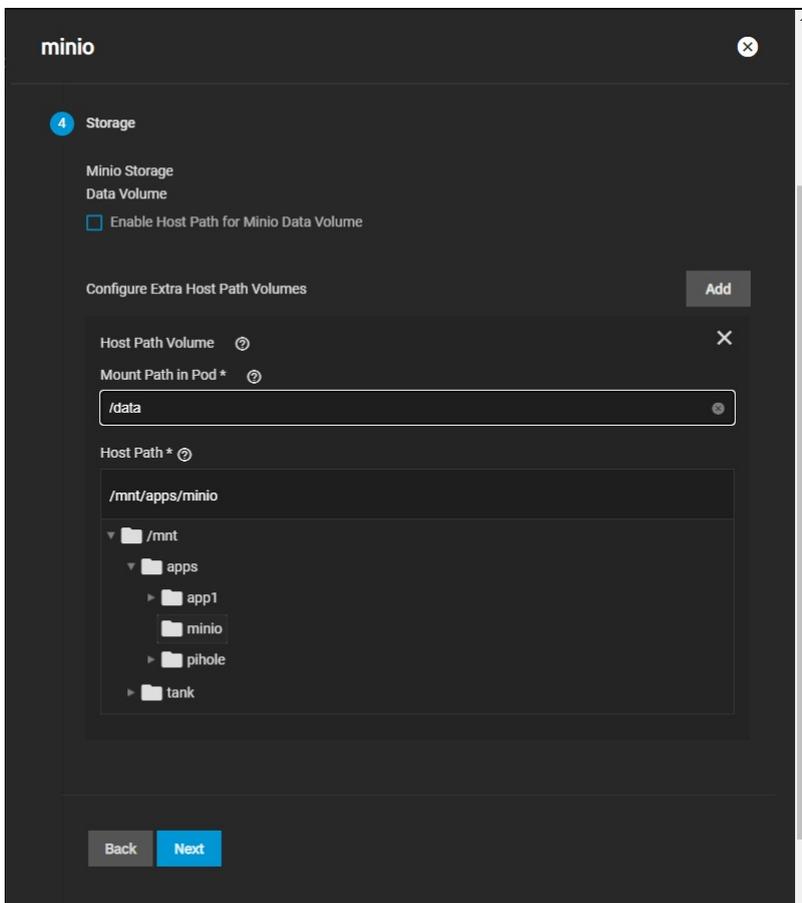


You can configure the API and UI access node ports and the MinIO domain name if you have TLS configured for MinIO. You can also configure a MinIO certificate if you wish.

Storage

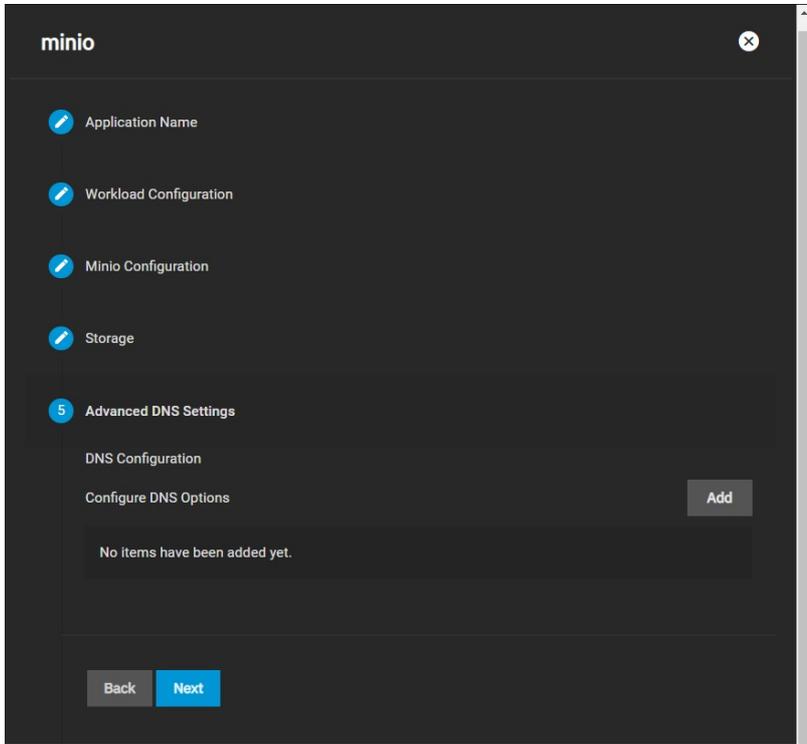
If you want to use a host path to store your MinIO data volume, select the **Enable Host Path for MinIO Data Volume** checkbox and select a path.

Under **Configure Extra Host Path Volumes**, enter the /data directory under **Mount Path in Pod**, then select the directory or dataset you created earlier and click **Next**.



Advanced DNS Settings

You can configure additional DNS options in Advanced DNS Settings. Click **Add** to add more DNS option entries. Click **Next**.



Confirm Options

Make sure the configuration summary meets your needs, then click **Save**.

11 - Reporting

- Reports Configuration
 - TrueCommand Enhancement

TrueNAS has a built-in reporting engine that provides helpful graphs and information about the system.

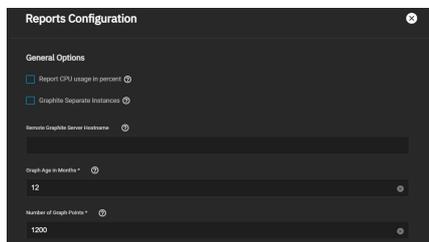


What does TrueNAS use for reporting?

TrueNAS uses [Graphite](#) to gather metrics and create visualizations.

The [Reports Configuration](#) menu lets you control how TrueNAS displays the graphs.

Reports Configuration



General Options

Name	Description
Report CPU usage in Percent	Reports CPU usage in percent instead of units of kernel time.
Graphite Separate Instances	Sends the <i>plugin instance</i> and <i>type instance</i> to Graphite as separate path components: <code>host.cpu.0.cpu.idle</code> . Disabling sends the <i>plugin</i> and <i>plugin instance</i> as one path component and <i>type</i> and <i>type instance</i> as another: <code>host.cpu-0.cpu-idle</code> .
Remote Graphite Server Hostname	Remote Graphite server Hostname or IP address.
Graph age in Months	Maximum time (in months) TrueNAS stores a graph (allowed values are 1-60). Changing this value causes the Confirm RRD Destroy dialog to appear. Changes do not take effect until TrueNAS destroys the existing reporting database.
Number of Graph Points	The number of points for each hourly, daily, weekly, monthly, or yearly graph (allowed values are 1-4096). Changing this value displays the Confirm RRD Destroy dialog. Changes do not take effect until TrueNAS destroys the existing reporting database.
Force	Forces TrueNAS to add the NTP server, even if it is unreachable.

TrueNAS clears the report history when you change *Report CPU*, *Graph Age*, or *Graph Points*.

TrueNAS saves and preserves reporting data across system upgrades and reboots so you can view usage trends over time.

Since TrueNAS writes reporting data frequently, do not store it on the boot pool or OS device. TrueNAS saves reporting data in `/var/db/collectd/rrd/`.

TrueCommand Enhancement

Want to increase TrueNAS's reporting functionality? Connect it to our TrueCommand multi-system management software.

TrueCommand [Reports](#) offer enhanced features like creating custom graphs and comparing utilization across multiple systems.

12 - System Settings

SCALE system management options are collected in this section of the UI and organized into a few different screens:

- **Update** controls when the system applies a new version. There are options to download and install an update, have the system check daily and stage updates, or apply a manual update file to the system.
- **General** shows system details and has basic, less intrusive management options, including web interface access, localization, and NTP server connections. This is also where users can input an Enterprise license or create a software bug ticket.
- **Advanced** contains options that are more central to the system configuration or meant for advanced users. Specific options include configuring the system console, log, and dataset pool, adding custom system controls, kernel-level settings, scheduled scripting or commands, and determining any isolated GPU devices. *Warning:* Advanced settings can be disruptive to system function if misconfigured.
- **Boot** lists each [ZFS](#) boot environment stored on the system. These restore the system to a previous version or specific point in time.
- **Services** displays each system component that runs continuously in the background. These typically control data sharing or other external access to the system. Individual services have their own configuration screens and activation toggles, and can be set to run automatically.
- **Shell** allows users to enter commands directly into the TrueNAS Operating System. Shell accepts Unix-like commands, and there is an experimental TrueNAS-specific command-line interface (CLI) for configuring the system separately from the web interface.
- **Enclosure** appears when the system is attached to compatible SCALE hardware. This is a visual representation of the system with additional details about disks and other physical hardware components.

Ready to get started? Choose a topic or article from the left-side Navigation pane. Click the < symbol to expand the menu to show the topics under this section.

12.1 - Update

- - [Automatic](#)
 - [Manual](#)

TrueNAS has several software branches (linear update paths) known as trains. SCALE is currently a Prerelease Train. Prerelease Trains have various preview/early build releases of the software.

SCALE has several trains available for updates. However, the web interface only displays trains you can select as an upgrade. For more information on other available trains, see [Truenas Upgrades](#).

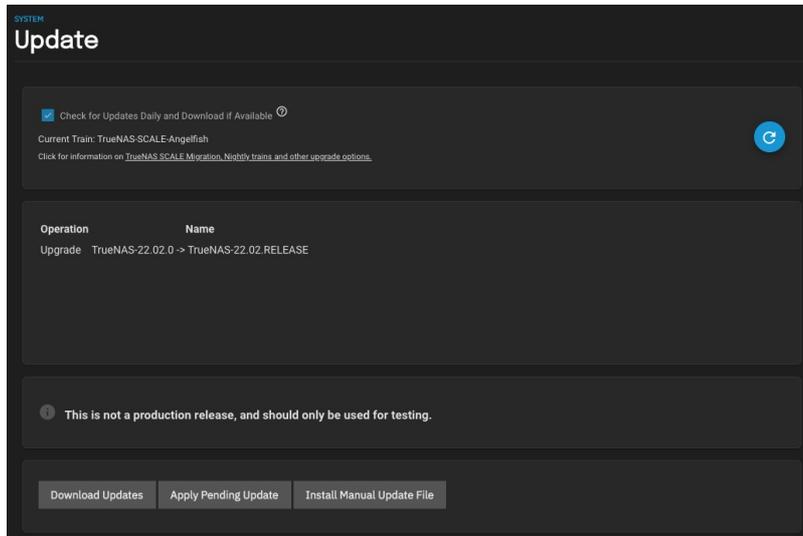
TrueNAS SCALE train is not suited for production use. Only use SCALE for testing.

Before using a non-production train, be prepared to experience bugs or problems. Testers are encouraged to submit bug reports at <https://jira.ixsystems.com>.

The TrueNAS SCALE Update screen lets users update their system using two different methods.

We recommend updating TrueNAS when the system is idle (no clients connected, no disk activity, etc). Most updates require a system reboot.

Update during scheduled maintenance times to avoid disrupting user activities.



Automatic

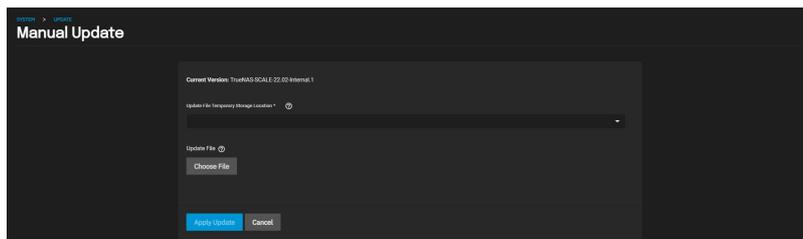
Set the *Check for Updates Daily and Download if Available* box to automatically download updates.

If an update is available, you can click *Apply Pending Update* to install it.

Manual

Download the [SCALE Manual Update File](#).

To manually update TrueNAS, click **Install Manual Update File** and save your configuration.

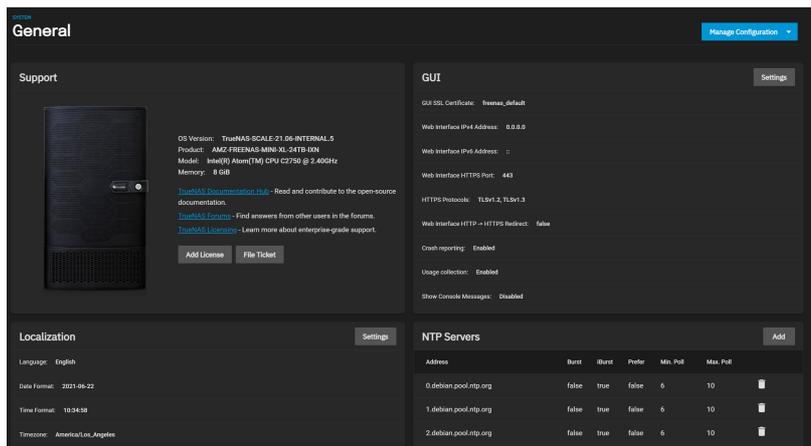


Select a temporary location to store the update file and click **Choose File**. Select the .iso you want to upgrade to and click **Apply Update**.

12.2 - General Settings

- [Manage Configuration](#)
- [Support](#)
- [GUI](#)
- [Localization](#)
- [NTP Servers](#)

The TrueNAS SCALE General Settings section provides settings options for support, graphic user interface, localization, NTP servers, and system configuration.



Manage Configuration

TrueNAS SCALE allows users to manage the system configuration via uploading/downloading configurations or resetting the system to the default configuration.

Download File

The *Download File* option downloads your TrueNAS SCALE's current configuration to the local machine.

When you download the configuration file, you have the option to *Export Password Secret Seed*, which includes encrypted passwords in the configuration file. This allows the configuration file to be restored to a different operating system device where the decryption seed is not already present. Users must physically secure configuration backups containing the seed to prevent unauthorized access or password decryption.

We recommend backing up the system configuration regularly. Doing so preserves settings when migrating, restoring, or fixing the system if it runs into any issues. Save the configuration file each time the system configuration changes.

Upload File

The *Upload File* option gives users the ability to replace the current system configuration with any previously saved TrueNAS SCALE configuration file.

All passwords will be reset if the uploaded configuration file was saved without the Password Secret Seed.

Reset to Defaults

The *Reset to Defaults* option resets the system's configuration to factory settings. After the configuration resets, the system will restart and users must set a new login password.

Save the system's current configuration with the *Download File* option before resetting the configuration to default settings.

If you do not save the system configuration before resetting it, you may lose data that was not backed up, and you will not be able to revert to the previous configuration.

Support

The *Support* window in the Advanced Settings screen displays the systems general hardware and software specs and contains links to the Documentation Hub, TrueNAS Forums, and enterprise licensing information.

There are also buttons that allow users to add an enterprise license or report bugs via a Jira support ticket.

GUI

The *GUI* window allows users to configure the TrueNAS SCALE web interface address.

Name	Description
GUI SSL Certificate	The system uses a self-signed certificate to enable encrypted web interface connections. To change the default certificate, select a different certificate that was created or imported in the Certificates section.
Web Interface IPv4 Address	Choose a recent IP address to limit the usage when accessing the administrative GUI. The built-in HTTP server binds to the wildcard address of 0.0.0.0 (any address) and issues an alert if the specified address becomes unavailable.
Web Interface IPv6 Address	Choose a recent IPv6 address to limit the usage when accessing the administrative GUI. The built-in HTTP server binds to the wildcard address of 0.0.0.0 (any address) and issues an alert if the specified address becomes unavailable.
Web Interface HTTPS Port	Allow configuring a non-standard port to access the GUI over HTTPS.
HTTPS Protocols	Cryptographic protocols for securing client/server connections. Select which Transport Layer Security (TLS) versions TrueNAS SCALE can use for connection security.
Web Interface HTTP -> HTTPS Redirect	Redirect HTTP connections to HTTPS. A GUI SSL Certificate is required for HTTPS. Activating this also sets the HTTP Strict Transport Security (HSTS) maximum age to 31536000 seconds (one year). This means that after a browser connects to the web interface for the first time, the browser continues to use HTTPS and renews this setting every year.
Crash Reporting	Send failed HTTP request data which can include client and server IP addresses, failed method call tracebacks, and middleware log file contents to iXsystems.
Usage Collection	Enable sending anonymous usage statistics to iXsystems.
Show Console Messages	Display console messages in real time at the bottom of the browser.

Localization

The *Localization* window lets users localize their system to a specific region.

Name	Description
Language	Select a language from the drop-down menu.
Date Format	Choose a date format.
Time Format	Choose a time format.
Timezone	Select a time zone.
Console Keyboard Map	Select a keyboard layout.

NTP Servers

The *NTP Servers* window allows user to configure Network Time Protocol (NTP) servers, which sync the local system time with an accurate external reference. By default, new installations use several existing NTP servers. TrueNAS SCALE supports adding custom NTP servers.

NTP Server Settings

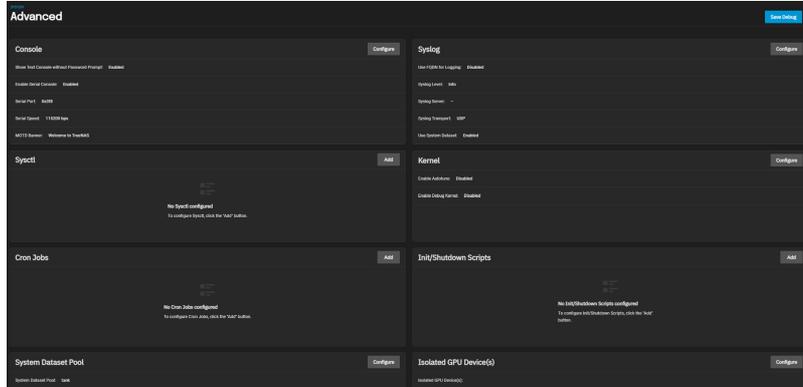
Name	Description
Address	Enter the hostname or IP address of the NTP server.
Burst	Recommended when Max. Poll is greater than 10. Only use on personal NTP servers or those under direct control. Do not enable when using public NTP servers.
IBurst	Speeds up the initial synchronization (seconds instead of minutes).
Prefer	Should only be used for highly accurate NTP servers such as those with time monitoring hardware.
Min Poll	The minimum polling interval, in seconds, as a power of 2. For example, 6 means 2^6 , or 64 seconds. The default is 6, minimum value is 4.
Max Poll	The maximum polling interval, in seconds, as a power of 2. For example, 10 means 2^{10} , or 1,024 seconds. The default is 10, maximum value is 17.
Force	Forces the addition of the NTP server, even if it is currently unreachable.

12.3 - Advanced Settings

- [Console](#)
 - [Syslog](#)
 - [Sysctl](#)
 - [Kernel](#)
 - [Cron Jobs](#)
 - [Init/Shutdown Scripts](#)
 - [System Dataset Pool](#)
 - [Isolated GPU Device\(s\)](#)
 - [Replication](#)

TrueNAS SCALE Advanced Settings provide configuration options for the Console, Syslog, Sysctl, Kernel, Cron Jobs, Init/Shutdown Scripts, System Dataset Pool, and Isolated GPU Device(s).

Advanced Settings have reasonable defaults in place. Changing advanced settings can be dangerous when done incorrectly. Please use caution before saving. Make sure you are comfortable with ZFS, Linux, and system [configuration backup and restoration](#) before making any changes.



Console

The **Console** window allows users to configure the [Console Setup menu](#).

Name	Description
Show Text Console without Password Prompt	Unset to add a login prompt to the system before showing the console menu.
Enable Serial Console	Do not set this if the Serial Port is disabled.
Serial Port	The serial console port address.
Serial Speed	The speed (in bits per second) the serial port uses.
MOTD Banner	The message that displays when a user logs in with SSH.

Syslog

The **Syslog** window allows users configure how and when the system sends log messages to the Syslog server.

Name	Description
Use FQDN for Logging	Set to include the Fully-Qualified Domain Name (FQDN) in logs to precisely identify systems with similar hostnames.
Syslog Level	When Syslog Server is defined, the system only sends logs matching this level.
Syslog Server	Remote Syslog server DNS hostname or IP address. You can use nonstandard port numbers by adding a colon and the port number to the hostname, like <code>mysyslogserver:1928</code> . Log entries are written to local logs and sent to the remote Syslog server.
Syslog Transport	Transport Protocol for the remote system log server connection. Choosing Transport Layer Security (TLS) also requires selecting a preconfigured system Certificate.
Use System Dataset	Store system logs on the system dataset. Unset to store system logs in <code>/var/</code> on the operating system device.

Sysctl

The **Sysctl** window allows users set up tunables that configure kernel parameters at runtime.

Name	Description
Variable	Enter the loader name, sysctl , or rc.conf variable to configure. TrueNAS uses loader tunables to specify parameters to pass to the kernel or load additional modules at boot time. rc.conf tunables enable system services and daemons and only take effect after a reboot. sysctl tunables configure kernel parameters while the system is running and generally take effect immediately.
Value	Enter a value to use for the loader, sysctl , or rc.conf variable.
Description	Enter a description of the tunable.
Enabled	Enable this tunable. Unset to disable this tunable without deleting it.

Kernel

The **Kernel** window contains options for system optimization and kernel debugging.

Name	Description
Enable Autotune	Activates a tuning script that attempts to optimize the system depending on the installed hardware. Warning: Autotuning is a temporary measure and is not a permanent fix for system hardware issues.
Enable Debug Kernel	Set to boot a debug kernel after the next system reboot.

Cron Jobs

The **Cron Jobs** window allows users to configure jobs that run specific commands or scripts on a regular schedule using [cron\(8\)](#). Cron Jobs help run repetitive tasks.

Name	Description
Description	Enter a description of the cron job.
Command	Enter the full path to the command or script to be run.

Run As User	Select a user account to run the command. The user must have permissions allowing them to run the command or script.
Schedule	Select a schedule preset or choose Custom to open the advanced scheduler. Note that an in-progress cron task postpones any later scheduled instance of the same task until the running task is complete.
Hide Standard Output	Hide standard output (stdout) from the command. When unset, TrueNAS mails any standard output to the user account cron that ran the command.
Hide Standard Error	Hide error output (stderr) from the command. When unset, TrueNAS mails any error output to the user account cron that ran the command.
Enabled	Enable this cron job. When unset, disable the cron job without deleting it.

Init/Shutdown Scripts

The **Init/Shutdown Scripts** window allows users to schedule commands or scripts to run at system startup or shutdown.

Name	Description
Description	Comments about this script.
Type	Select Command for an executable or Script for an executable script.
Command	Enter the command with any options.
Script	Select the script. The script will be run using dash(1) .
When	Select when the command or script runs:
<i>Pre Init</i> is early in the boot process, after mounting filesystems and starting networking.	
<i>Post Init</i> is at the end of the boot process, before FreeNAS services start.	
<i>Shutdown</i> is during the system power-off process.	
Enabled	Enable this cron job. When unset, disable the cron job without deleting it.
Timeout	Automatically stop the script or command after the specified seconds.

System Dataset Pool

System Dataset Pool allows users select the storage pool to hold the system dataset. The system dataset stores debugging core files, encryption keys for encrypted pools, and Samba4 metadata such as the user and group cache and share level permissions.

Users can move the system dataset to unencrypted pools or encrypted pools without passphrases.

Users can move the system dataset to a key encrypted pool but, after the move, the pool encryption type can't be changed to passphrase. If the encrypted pool already has a passphrase set, you cannot move the system dataset to that pool.

Isolated GPU Device(s)

The **Isolated GPU Device(s)** window allows users to isolate additional GPU devices for GPU passthrough.

GPU passthrough allows the TrueNAS SCALE kernel to directly present an internal PCI GPU to a virtual machine (VM).

The GPU device acts like the VM is driving it, and the VM detects the GPU as if it is physically connected.

Replication

The **Replication** window allows users to limit the maximum number of replication tasks executed simultaneously.

12.4 - Boot Environments

- [Managing Boot Environments](#)
- [Boot Actions](#)
- [Changing Boot Environments](#)

TrueNAS supports a ZFS feature known as boot environments. These are snapshot clones that TrueNAS can boot into. Only one boot environment can be used for booting.

How does this help me?

A boot environment allows rebooting into a specific point in time and greatly simplifies recovering from system misconfigurations or other potential system failures. With multiple boot environments, the process of updating the operating system becomes a low-risk operation. The updater automatically creates a snapshot of the current boot environment and adds it to the boot menu before applying the update. If anything goes wrong during the update, the system administrator can boot TrueNAS into the previous environment to restore system functionality.

Managing Boot Environments

To view the list of boot environments on the system, go to **System Settings > Boot**. Each boot environment entry contains this information:

- **Name**: the name of the boot entry as it appears in the boot menu.
- **Active**: indicates which entry boots by default if a boot environment is not active.
- **Created**: indicates the boot environment creation date and time.
- **Space**: shows boot environment size.
- **Keep**: indicates whether or not TrueNAS deletes this boot environment when a system update does not have enough space to proceed.

To access more options for a boot environment, click **⋮**:

Activate

Only appears on entries which are not currently set to **Active**. Activating an environment means the system boots into the point of time saved in that environment the next time it is started. The status changes to **Reboot** and the current **Active** entry changes from **Now/Reboot** to **Now**, indicating that it is the currently booted environment but will not be used on next boot.

Clone

Copy the selected boot environment into a new entry. The clone *Name* only allows alphanumeric characters, dashes (-), underscores (_), and periods (.) are allowed.

Rename

Changes the boot environment name. Alphanumeric characters, dashes (-), underscores (_), and periods (.) are allowed.

Delete

Removes the highlighted entry and also removes that entry from the boot menu. The **default** and any **Active** entries cannot be deleted. Because an activated entry cannot be deleted, this button does not appear for the active boot environment. To delete a currently **Active** entry, first activate another entry.

Keep

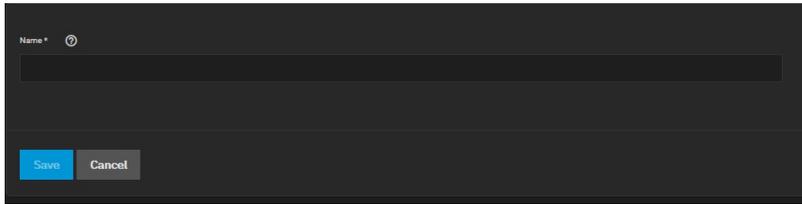
Toggles whether or not the updater can prune (automatically delete) this boot environment if there is not enough space to proceed with the update.

Boot Actions

Click **ACTIONS** to:

Add

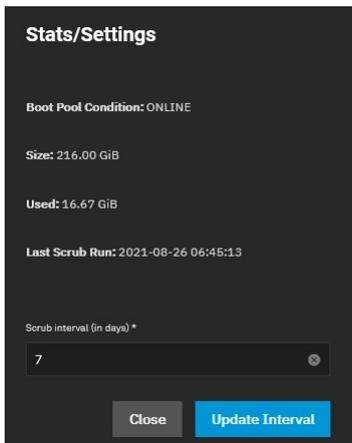
Make a new boot environment from the active environment:



Only alphanumeric characters, dashes (-), and underscores (_) are allowed in the *Name*. *Name* the new boot environment and click *Save*.

Stats/Settings

Display statistics for the operating system device: **Boot pool Condition**, **Size** and **Used**, and **Last Scrub Run**. By default, the operating system device is scrubbed every 7 days. To change the default, input a different number in the *Scrub interval (in days)* field and click *Update Interval*.



Boot Pool Status

Shows the status of each device in the operating system device (boot-pool), including any read, write, or checksum errors.



Name	Read	Write	Checksum	Status
bootpool	0	0	0	OK
bootpool2	0	0	0	OK

Scrub Boot Pool

Perform a manual “scrub” (data integrity check) of the operating system device.

Changing Boot Environments

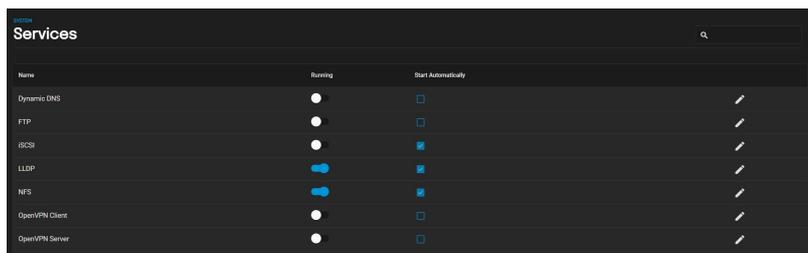
Sometimes, rolling back to an older boot environment can be useful. For example, if an update process doesn't go as planned, it is easy to roll back to a previous boot environment. TrueNAS automatically creates a boot environment when the system updates.

To activate a different boot environment, go to **System Settings > Boot** and click **Activate** for the desired boot environment. Next, click **Activate**. This boot environment shows **Reboot** in the **Active** column. This means the boot environment becomes active on the next system boot. The system configuration also changes to the state it was in when the boot environment was created.

12.5 - Services

System Settings > Services displays each system component that runs continuously in the background. These typically control data-sharing or other external access to the system. Individual services have configuration screens and activation toggles, and you can set them to run automatically.

Documented services related to data sharing or automated tasks are in their respective [Shares](#) and [Tasks](#) articles.

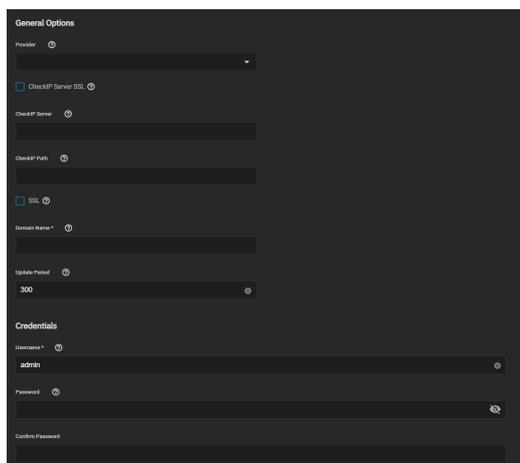


Dynamic DNS

[Dynamic Domain Name Service \(DDNS\)](#) is useful when you connect TrueNAS to an Internet service provider (ISP) that periodically changes the system's IP address. With Dynamic DNS, the system automatically associates its current IP address with a domain name and continues to provide access to TrueNAS even if the system IP address changes.

Configuring Dynamic DNS

DDNS requires registration with a DDNS service such as [DynDNS](#) before configuring TrueNAS. Have the DDNS service settings available or open in another browser tab when configuring TrueNAS. Log in to the TrueNAS web interface and go to **Services > Dynamic DNS**.



General Options

Name	Description
Provider	Several providers are supported. If a specific provider is not listed, select Custom Provider and enter the information in the Custom Server and Custom Path fields.
CheckIP-Server SSL	Use HTTPS for the connection to the CheckIP Server.
CheckIP Server	Name and port of the server that reports the external IP address. For example, entering checkip.dyndns.org:80 uses Dyn IP detection to discover the remote socket IP address.
CheckIP Path	Path to the CheckIP Server. For example, no-ip.com uses a CheckIP Server of dynamic.zoneedit.com and CheckIP Path of /checkip.html.
SSL	Use HTTPS for the connection to the server that updates the DNS record.
Domain Name	Fully qualified domain name of the host with the dynamic IP address. Separate multiple domains with a space, comma (,), or semicolon (;). Example: myname.dyndns.org; myothername.dyndns.org.
Update Period	How often the IP is checked in seconds.

Credentials

Name	Description
Username	Username for logging in to the provider and updating the record.
Password	Password for logging in to the provider and updating the record.

Your DDNS solution provides the required values for the fields. Start the DDNS service after choosing your **Provider** options and saving the settings.

FTP, SFTP, and TFTP

The [File Transfer Protocol \(FTP\)](#) is a simple option for data transfers. The SSH and Trivial FTP options provide secure or simple config file transfer methods respectively.

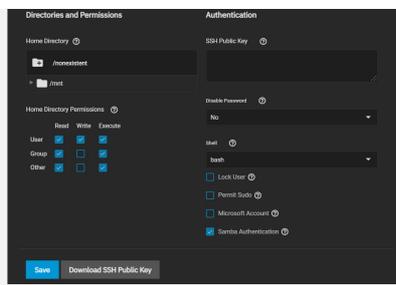
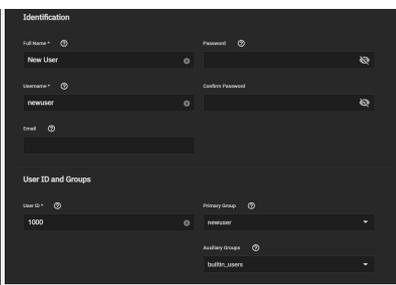
Options for configuring **FTP**, **SSH**, and **TFTP** are in **System Settings > Services**. Click the [FTP](#) icon to configure the related service.

FTP

FTP requires a new dataset and a local user account.

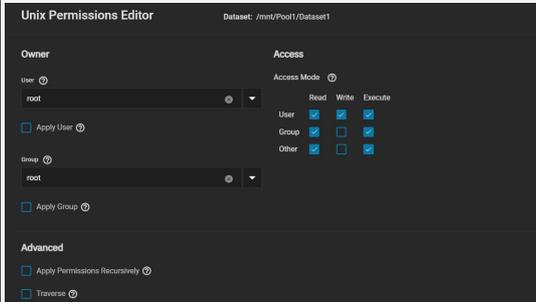
Go to **Storage** to add a new [dataset](#).

Next, go to **Credentials > Local Users** and click **Add** to create a local user on the TrueNAS.



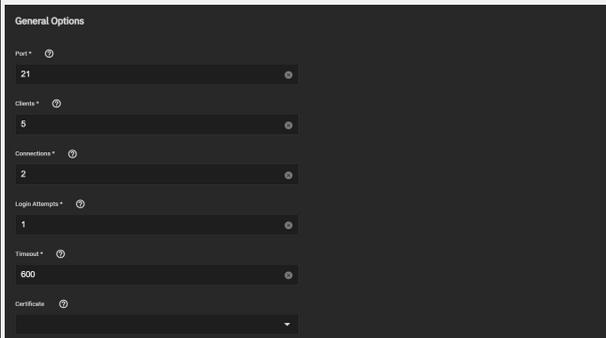
Assign a user name and password, and link the newly created FTP share dataset as the user home directory. You can do this for every user, or create a global account for FTP (for example, *OurOrgFTPacnt*).

Return to **Storage**, find the new dataset, click **View Permissions**, and select **View Permissions**. Next click **Apply User** and **Apply Group** before saving.



Service Configuration

To configure FTP, go to **System Settings > Services** and find **FTP**, then click **Configure**.



Configure the options according to your environment and security considerations.

General Options

Name	Description
Port	Set the port the FTP service listens on.
Clients	The maximum number of simultaneous clients.
Connections	Set the maximum number of connections per IP address. 0 is unlimited.
Login Attempts	Enter the maximum attempts before client is disconnected. Increase if users are prone to typos.
Timeout	Maximum client idle time in seconds before disconnect.
Certificate	The SSL certificate to be used for TLS FTP connections. To create a certificate, go to Certificates .

Advanced

Access

Name	Description
Always Chroot	Set to only let users access their home directory if they are in the wheel group. This option increases security risk.
Allow Root Login	Allow anonymous FTP logins with access to the directory specified in <i>Path</i> .
Allow Anonymous Login	Allow any local user to log in. By default, only members of the <i>ftp</i> group are allowed to log in.
Allow Local User Login	Setting this option results in timeouts when <i>identd</i> is not running on the client.
Require IDENT Authentication	Sets default permissions for newly created files.
File Permissions	Sets default permissions for newly created directories.

TLS

Name	Description
Enable TLS	Allow encrypted connections. Requires a certificate (created or imported in Certificates).
TLS Policy	Define whether the control channel, data channel, both channels, or neither channel of an FTP session must occur over SSL/TLS. The policies are described here .
TLS Allow Client Renegotiations	We don't recommend this, since it breaks security measures. See mod_tls for details.
TLS Allow Dot Login	If set, TrueNAS checks the user home directory for a <i>.tlogin</i> file containing one or more PEM-encoded certificates. If not found, the user is prompted for password authentication.
TLS Allow Per User	If set, allows user password to be sent unencrypted.
TLS Common Name Required	When set, the common name in the certificate must match the FQDN of the host.
TLS Enable Diagnostics	If set when troubleshooting a connection, logs more verbosely.
TLS Export Certificate Data	Set to export the certificate environment variables.

TLS No Certificate Request	Set if the client cannot connect from poorly handling the server certificate request.
TLS No Empty Fragments	We don't recommend this option, since it bypasses a security mechanism.
TLS No Session Reuse Required	This option reduces connection security. Only use it if the client does not understand reused SSL sessions.
TLS Export Standard Vars	If selected, sets several environment variables.
TLS DNS Name Required	If set, the client DNS name must resolve to its IP address and the cert must contain the same DNS name.
TLS IP Address Required	If set, the client certificate IP address must match the client IP address.

Bandwidth

Name	Description
Local User Upload Bandwidth: (Examples: 500 KiB, 500M, 2 TB) *	This field accepts human-readable input in KiBs or greater (M, GiB, TB, etc.). Default 0 KiB is unlimited.
Local User Download Bandwidth	This field accepts human-readable input in KiBs or greater (M, GiB, TB, etc.). Default 0 KiB is unlimited.
Anonymous User Upload Bandwidth	This field accepts human-readable input in KiBs or greater (M, GiB, TB, etc.). Default 0 KiB is unlimited.
Anonymous User Download Bandwidth	This field accepts human-readable input in KiBs or greater (M, GiB, TB, etc.). Default 0 KiB is unlimited.

Other Options

Name	Description
Minimum Passive Port	Used by clients in PASV mode. A default of 0 means any port above 1023.
Maximum Passive Port	Used by clients in PASV mode. A default of 0 means any port above 1023.
Enable FXP	Enable File eXchange Protocol. We don't recommend this, since it leaves the server vulnerable to FTP bounce attacks.
Allow Transfer Resumption	Set to allow FTP clients to resume interrupted transfers.
Perform Reverse DNS Lookups	Performs reverse DNS lookups on client IPs. Causes long delays if reverse DNS isn't configured.
Masquerade Address	Public IP address or hostname. Set if FTP clients cannot connect through a NAT device.
Display Login	The message shown to local login users after authentication. Not shown to anonymous login users.
Auxiliary Parameters	Used to add additional proftpd(8) parameters.

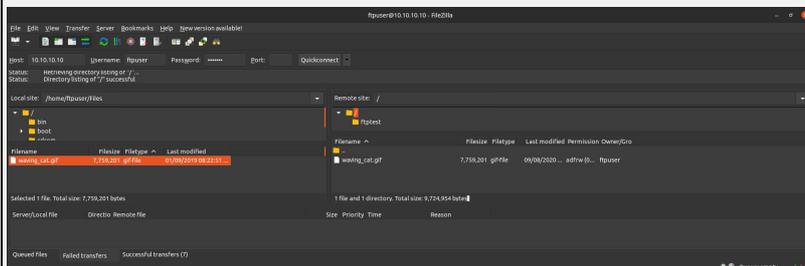
To confine FTP sessions to a local user's home directory, ensure **chroot** is enabled and allow **Local User Login**.

Do *not* allow anonymous or root access unless it is necessary. For better security, enable TLS when possible (especially when exposing FTP to a WAN). TLS effectively makes this [FTPS](#).

FTP Connection

Use a browser or FTP client to connect to the TrueNAS FTP share. The images below use [FileZilla](#), a free option.

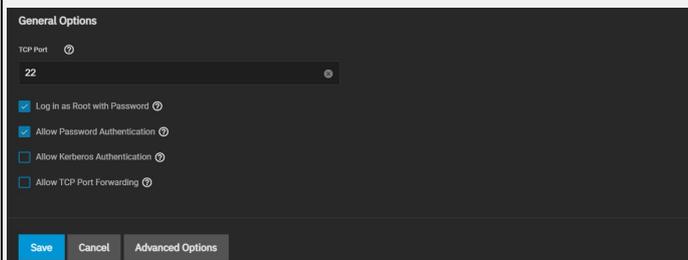
The user name and password are those of the local user account on the TrueNAS. The default directory is the same as the user's /home directory. After connecting, you can create directories and upload/download files.



SFTP

SFTP (SSH File Transfer Protocol) is available by enabling SSH remote access to the TrueNAS system. SFTP is more secure than standard FTP as it applies SSL encryption on all transfers by default.

Go to **System Settings > Services**, find the **SSH** entry, and click the



Set **Allow Password Authentication** and decide if you need **Log in as Root with Password**.

SSH with root is a security vulnerability. It allows users to fully control the NAS remotely with a terminal instead of providing SFTP transfer access.

Review the remaining options and configure them according to your environment or security needs.

General Options

Name	Description
TCP Port	Open a port for SSH connection requests.
Log in as Root with Password	Root logins are discouraged. Allows root logins. A password must be set for the root user account.
Allow Password	Enabling allows SSH login authentication using a password. Warning: when directory services are enabled, this setting grants access to all users the directory service imported. When disabled, authentication requires keys for all users (requires additional SSH client and server

Authentication	setup).
Allow Kerberos Authentication	Before enabling, ensure valid entries exist in Directory Services (Kerberos Realms and Keytabs) and the system can communicate with the Kerberos Domain Controller.
Allow TCP Port Forwarding	Set to let users bypass firewall restrictions using the SSH port forwarding feature .

Advanced Options

Name	Description
Bind Interfaces	Select interfaces for SSH to listen on. Leave all options unselected for SSH to listen on all interfaces.
Compress Connections	Select the syslog(3) level of the SFTP server.
SFTP Log Level	Select the syslog(3) facility of the SFTP server.
SFTP Log Facility	Allow more ciphers for sshd(8) in addition to the defaults in sshd_config(5) . <i>None</i> allows unencrypted SSH connections and AES128-CBC allows the 128-bit Advanced Encryption Standard .
Weak Ciphers	WARNING: these ciphers are security vulnerabilities. Only allow them in a secure network environment.
Auxiliary Parameters	Add any more sshd_config(5) options not covered in this screen. Enter one option per line. These options are case-sensitive. Typos can prevent the SSH service from starting.

SFTP Connections

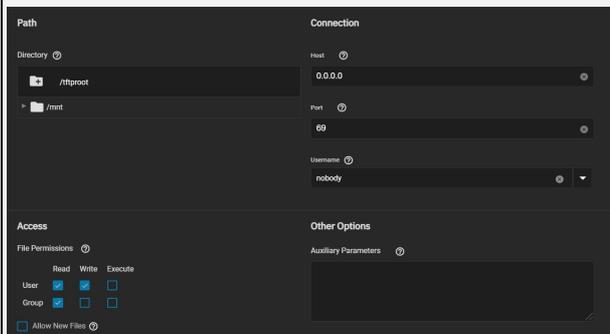
Open an FTP client (like FileZilla) or command line. This article shows using FileZilla as an example. Using FileZilla, enter *SFTP://TrueNAS IP*, *'username'*, *'password'*, and port **22** to connect.

SFTP does not offer chroot locking. While chroot is not 100% secure, lacking chroot lets users move up to the root directory and view internal system information. If this level of access is a concern, FTP with TLS might be the more secure choice.

TFTP

The Trivial File Transfer Protocol (TFTP) is a lightweight version of FTP typically used to transfer configuration or boot files between machines, such as routers, in a local environment. TFTP provides a limited set of commands and provides no authentication.

If TrueNAS is only storing images and configuration files for network devices, configure and start the TFTP service. Starting the TFTP service opens UDP port **69**.



Path

Name	Description
Directory	Browse to an existing directory to use for storage. Some devices can require a specific directory name. Consult the documentation for that device to see if there are any restrictions.

Connection

Name	Description
Host	The default host to use for TFTP transfers. Enter an IP address. Example: 192.0.2.1
Port	The UDP port number that listens for TFTP requests. Example: 8050
Username	Select the account to use for TFTP requests. This account must have permission to the Directory.

Access

Name	Description
File Permissions	Adjust the file permissions using the checkboxes.
Allow New Files	Set when network devices need to send files to the system.

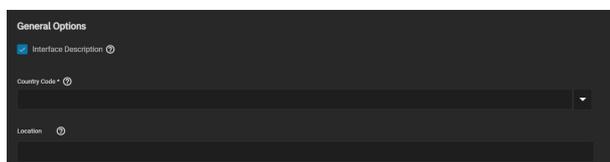
Other Options

Name	Description
Auxiliary Parameters	Add more options from tftpd . Add one option on each line.

LLDP

Network devices use the [Link Layer Discovery Protocol \(LLDP\)](#) to advertise their identity, capabilities, and neighbors on an Ethernet network. TrueNAS uses the [ladvd](#) LLDP implementation. When the local network contains managed switches, configuring and starting LLDP tells TrueNAS to advertise itself on the network.

To configure LLDP, go to **System Settings > Services**, find **LLDP** and click the



General Options

Name	Description
Interface Description	Enables receive mode. Any received peer information is saved in interface descriptions.

County Code	Two-letter ISO 3166-1 alpha-2 code used to enable LLDP location support.
Location	The physical location of the host.

Set **Interface Description** and enter a **Country Code** before enabling the LLDP service.

OpenVPN

A virtual private network (VPN) is an extension of a private network over public resources. It lets clients securely connect to a private network even when remotely using a public network. TrueNAS provides [OpenVPN](#) as a system-level service to provide VPN server or client functionality. TrueNAS can act as a primary VPN server that allows remote clients to access system data using a single TCP or UDP port. Alternatively, TrueNAS can integrate into a private network, even when the system is in a separate physical location or only has access to publicly visible networks.

Before configuring TrueNAS as either an OpenVPN server or client, you need an existing public key infrastructure (PKI) with [Certificates and Certificate Authorities](#) created in or imported to TrueNAS.

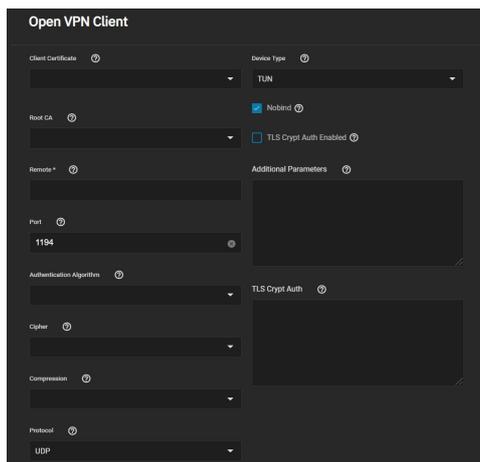
What does this do?

Certificates allow TrueNAS to authenticate with clients or servers by confirming a valid master Certificate Authority (CA) signed the network credentials. To read more about the required PKI for OpenVPN, see the [OpenVPN PKI Overview](#).

In general, configuring TrueNAS OpenVPN (server or client) includes selecting networking credentials, setting connection details, and choosing additional security or protocol options.

OpenVPN Client

Go to **System Settings > Services** and find **OpenVPN Client**. Click the  to configure the service.



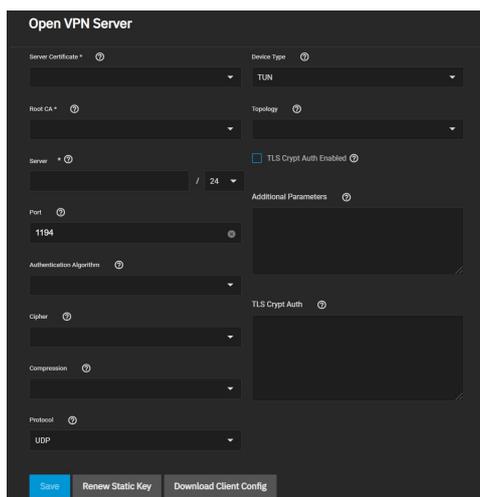
Choose the certificate to use as an OpenVPN client. The certificate must exist in TrueNAS and be active (unrevoked). Enter the **Remote** OpenVPN server's hostname or IP address.

Continue to review and choose any other [Connection Settings](#) that fit your network environment and performance requirements. The **Device Type** must match the OpenVPN server **Device Type**. **Nobind** prevents using a fixed port for the client and is enabled by default so the OpenVPN client and server run concurrently.

Finally, review the [Security Options](#) and ensure they meet your network security requirements. If the OpenVPN server uses TLS Encryption, copy the static encryption key and paste it into the **TLS Crypt Auth** field.

OpenVPN Server

Go to **System Settings > Services** and find **OpenVPN Server**. Click the  to configure the service.



Choose a **Server Certificate** for the OpenVPN server. The certificate must exist in TrueNAS and be active (unrevoked).

Now define an IP address and netmask for the OpenVPN **Server**. Select the remaining [Connection Settings](#) that fit your network environment and performance requirements. If using a **TUN Device Type**, you can choose a virtual addressing topology for the server in **Topology**:

- **NET30**: Use one /30 subnet per client in a point-to-point topology. Use when connecting clients are Windows systems.
- **P2P**: Point-to-point topology that points the local server and remote client endpoints to each other. Each client gets one IP address. Use when none of the clients are Windows systems.
- **SUBNET**: The interface uses an IP address and subnet. Each client gets one IP address. Windows clients require the **TAP-Win32 driver** version 8.2 or newer. **TAP** devices always use the **SUBNET Topology**.

TrueNAS applies the **Topology** selection to any connected clients.

When **TLS Crypt Auth Enabled** is selected, TrueNAS generates a static key for the **TLS Crypt Auth** field after saving the options. To change this key, click **Renew Static Key**. Clients connecting to the server require the key. TrueNAS stores keys in the system database and includes them in client config files. We recommend always backing up keys in a secure location.

Finally, review the [Security Options](#) and choose settings that meet your network security requirements.

After configuring and saving your OpenVPN Server, generate client configuration files to import to any OpenVPN client systems connecting to this server. You need the certificate from the client system already imported into TrueNAS. To generate the configuration file, click **Download Client Config** and select the

Client Certificate.

Common Options (Client or Server)

Many OpenVPN server or client configuration fields are identical. This section covers these fields and lists specific configuration options in the [Server](#) and [Client](#) sections.

The **Additional Parameters** field manually sets any core OpenVPN config file options. See the OpenVPN [Reference Manual](#) for descriptions of each option.

Connection Settings

Setting	Description
Root CA	The Certificate Authority (CA) must be the root CA you used to sign the client and server certificates.
Port	The port that the OpenVPN connection is to use.
Compression	Choose a compression algorithm for traffic. Leave empty to send data uncompressed. LZO is a standard compression algorithm that is backward compatible with previous (pre-2.4) versions of OpenVPN. LZ4 is newer and typically faster and requires fewer system resources.
Protocol	Choose between UDP or TCP OpenVPN protocols. UDP sends packets in a continuous stream. TCP sends packets sequentially. UDP is usually faster and less strict about dropped packets than TCP. To force the connection to be IPv4 or IPv6, choose one of the 4 or 6 UDP or TCP options.
Device Type	Use a TUN or TAP virtual networking device and layer with OpenVPN. The device must be identical between the OpenVPN server and clients.

Security Options

OpenVPN includes several security options since using a VPN involves connecting to a private network while sending data over less secure public resources. Security options are not required, but they help protect data users send over the private network.

Setting	Description
Authentication Algorithm	Validates packets sent over the network connection. Your network environment might require a specific algorithm. If not, SHA1 HMAC is a reliable algorithm to use.
Cipher	Encrypts data packets sent through the connection. Ciphers aren't required but can increase connection security. You might need to verify which ciphers your networking environment requires. If there are no specific cipher requirements, AES-256-GCM is a good default choice.
TLS Encryption	When TLS Crypt Auth Enabled is selected, OpenVPN adds another layer of security by encrypting all TLS handshake messages. This setting requires sharing a static key between the OpenVPN server and clients.

Service Activation

Click **Save** after configuring the server or client service. Start the service by clicking the related toggle in **System Settings > Services**. Hover over the toggle to check the service current state.

Selecting **Start Automatically** starts the service whenever TrueNAS completes booting.

S3

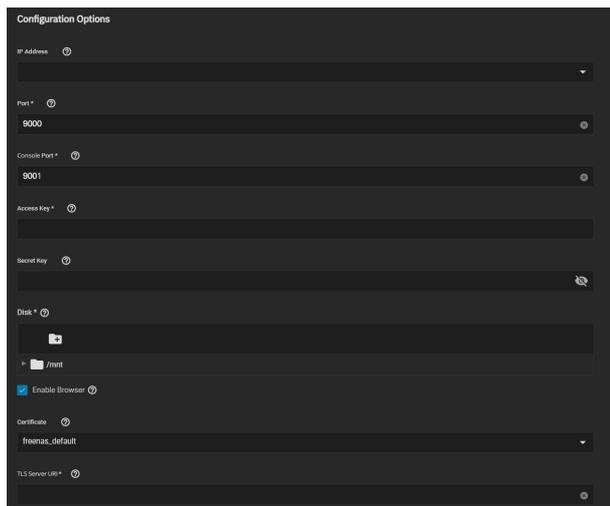
S3 allows you to connect to TrueNAS from a networked client system with the Minio browser, s3cmd, or S3 browser.

Background

S3 is an object storage protocol that many major cloud providers like Amazon Web Services™ use. On TrueNAS, the service is another way to store files and can be viewed with a web browser. Because S3 is the de facto standard for cloud-based storage, setting up an S3 service allows organizations or online application developers to use TrueNAS to replace or archive expensive cloud storage.

Setting up the S3 service

Go to the **System Settings > Services** and find **S3**, then click [Configure](#) to configure the service.



Field Descriptions

S3 Configuration Options

Name	Description
IP Address	Enter the IP address that runs the S3 service. <i>0.0.0.0</i> tells the server to listen on all IPv4 addresses. <i>::</i> allows the same for any IPv6 address. Select the TrueNAS IP address to constrain it to a specific network.
Port	Enter a static port for the MinIO web console. Default is 9001.
Console Port	Enter the TCP port that provides the S3 service.
Access Key	Enter the S3 access ID. See Access keys for more information.
Secret Key	Enter the S3 secret access key. See Access keys for more information.
Disk	Browse to a directory to define the S3 filesystem path.

Enable Browser	Enables the S3 service web UI. Access the MinIO web UI by entering the IP address and port number separated by a colon in the browser address bar. Example: 192.168.1.0:9000.
Certificate	Use an SSL certificate created or imported in Credentials > Certificates for secure S3 connections.
TLS Server URI	If using an SSL certificate, enter the MinIO server's proxy-able address

Select a clean dataset, one that doesn't have existing data files. Minio manages files as objects that you *cannot* mix with other dataset files. You can create new datasets by going to **Storage** and clicking **> Add Dataset.**

Configure the remaining options as needed in your environment and start the service after saving any changes.

Minio Connections

When **Enable Browser** is selected, test Minio browser access by opening a web browser and typing the TrueNAS IP address with the TCP port. You must allow the chosen **Port** through the network firewall to permit creating buckets and uploading files. Example: <https://192.168.0.3:9000>.

Minio supports two different connection methods.

s3cmd

Linux or macOS users must have the [s3cmd](#) service installed before beginning this setup. On Windows, users can also refer to [S3Express](#) for a similar command-line experience.

Ubuntu or other Linux distributions can access the configuration by running `s3cmd --configure` to walk through critical settings.

Enter the specified access key and the secret key. Under the **S3 Endpoint**, enter the TrueNAS IP address followed by TCP port, and reply **N** to the DNS-style bucket+hostname.

Save the file. On Linux, the default is in the home directory `~/.s3cfg`.

If the connection has any issues, open `.s3cfg` again to troubleshoot. In Ubuntu, use `nano .s3cfg` or `vi .s3cfg` or `gedit .s3cfg` depending on the preferred text editor. For other operating systems, `.s3cfg` file location and editing tools might vary.

Scroll down to the `host_bucket` area and ensure the configuration removed the `%(bucket)s.` portion and the address points to the `IP_address:TCP_port` for the system.

Correct Example

```
host_base = `192.168.123.207:9000`
host_bucket = `192.168.123.207:9000`
```

Incorrect Example

```
host_base = `192.168.123.207`
host_bucket = `%(bucket)s.192.168.123.207`
```

Poll the buckets using `s3cmd ls` to see the buckets created with the Minio browser.

For more information on using Minio with `s3cmd`, see <https://docs.minio.io/docs/s3cmd-with-minio.html> and <https://s3tools.org/s3cmd>.

S3 Browser (Windows)

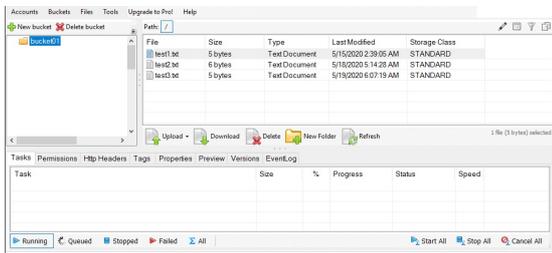
The Windows PC S3 browser is another convenient way to connect to the Minio S3 from TrueNAS.

To set it up, first [install the S3 browser](#).

After installation completes, add a new account.

In the settings, select **S3 Compatible Storage** as the **Account Type**, then enter the Minio access point similar to the `s3cmd` setup (TrueNAS_IP_address:9000 or other port if set differently). Select the SSL settings appropriate for the particular setup. The S3 browser assumes SSL by default, but it can be unset for a LAN attached session.

It is possible to access, create new buckets, or upload files to created buckets.



SNMP

[SNMP \(Simple Network Management Protocol\)](#) monitors network-attached devices for conditions that warrant administrative attention. TrueNAS uses [Net-SNMP](#) to provide SNMP. To configure SNMP, go to **System Settings > Services** page, find **SNMP**, and click the

Field Descriptions

General Options

Name	Description
Location	Enter the location of the system.
Contact	E-mail address that receives SNMP service messages.
Community	Change from public to increase system security. Can only contain alphanumeric characters, underscores (<code>_</code>), dashes (<code>-</code>), periods (<code>.</code>), and spaces. This can be left empty for SNMPv3 networks.

SNMP v3 Options

Name	Description
SNMP v3 Support	Set to to enable support for SNMP version 3 . See snmpd.conf(5) for configuration details.
Username	Enter a username to register with this service.
Authentication Type	Choose an authentication method: --- for none, SHA , or MD5
Password	Enter a password of at least eight characters.
Privacy Protocol	Choose a privacy protocol: --- for none, AES , or DES
Privacy Passphrase	Enter a separate privacy passphrase. If field is left blank, the default <i>Password</i> is used.

Other Options

Name	Description
Auxiliary Parameters	Enter any additional snmpd.conf options. Add one option for each line.
Expose zillstat via SNMP	Enabling this option may have performance implications on your pools.
Log Level	Choose how many log entries to create. Choices range from least (Emergency) to most (Debug).

Port **UDP 161** listens for SNMP requests when starting the SNMP service.

Management Information Bases (MIBs)

Available Management Information Bases (MIBs) are located in `/usr/local/share/snmp/mibs`. This directory contains many files routinely added or removed from the directory. Check the directory on your system by going to **System Settings > Shell** and entering `ls /usr/local/share/snmp/mibs`. Here is a sample of the directory contents:

```
Linux truenas.ixsystems.com 5.10.42#truenas #1 SMP Mon Aug 30 21:54:59 UTC 2021 x86_64

TrueNAS (c) 2009-2021, iXsystems, Inc.
All rights reserved.
TrueNAS code is released under the modified BSD license with some
files copyrighted by (c) iXsystems, Inc.

For more information, documentation, help or support, go here:
http://truenas.com
Welcome to TrueNAS
Last login: Mon Sep 27 09:12:39 PDT 2021 from 10.231.1.215 on pts/0
truenas# ls /usr/local/share/snmp/mibs
FREEMIB.txt  IRI-SENSORS-MIB.txt
truenas#
```

SSH



Video Player is loading.
 Video URL: <https://www.truenas.com/docs/files/scaleangelfishsshaccess.mp4>

Play Video

The SSH service lets users connect to TrueNAS with the [Secure Shell Transport Layer Protocol](#). When using TrueNAS as an SSH server, the users in the network must use [SSH client software](#) to transfer files with SSH.

Current Time 0:00

Allowing external connections to TrueNAS is a security vulnerability! Do not enable SSH unless you require external connections.

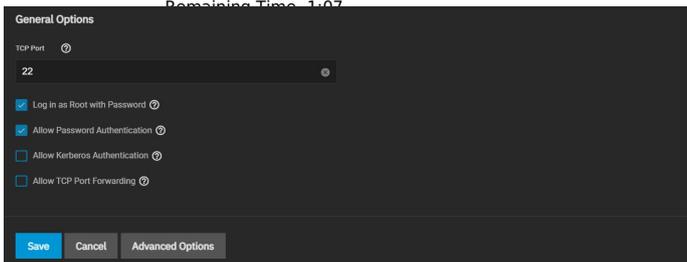
Duration 1:07

Activate or configure the SSH service on the **System Settings > Services** page.

Stream Type LIVE

To configure SSH go to **System Settings > Services**, find **SSH**, and click

Remaining Time: 1:07



Captions

Configure the options as needed to match your network environment.

SSH Service Fields	
• captions off, selected	
General Options Audio Track	
Name	Description
TCP Port	Open a port for SSH connection requests.
Log in as Root with Password	Root logins are discouraged. Allows root logins. A password must be set for the root user account.
Allow Password Authentication	Enabling allows SSH login with the system using a password. When directory services are enabled, this setting grants access to all users the directory service imported. When disabled, authentication requires keys for all users (requires additional SSH client and server setup).
Allow Kerberos Authentication	Before enabling, ensure valid entries exist in Directory Services (Kerberos Realms and Keytabs) and the system can communicate with the Kerberos Domain Controller.
Allow TCP Port Forwarding	Set to let users bypass firewall restrictions using the SSH port forwarding feature .
Advanced Options Color Black Transparency Transparent	
Name	Description
Bind Interfaces	Select interfaces for SSH to listen on. Leave all options unselected for SSH to listen on all interfaces.
Compress	Select the syslog(3) level of the SFTP server.

Connections	Text Edge Style
SFTP Log Level	Select the syslog(3) facility of the SFTP server.
SFTP Log Facility	Allow more ciphers for sshd(8) in addition to the defaults in sshd_config(5) . <i>None</i> allows unencrypted SSH connections and AES128-CBC allows the 128-bit Advanced Encryption Standard . <small>Proportional Sans-Serif</small>
Weak Ciphers	WARNING: these ciphers are security vulnerabilities. Only allow them in a secure network environment.
Auxiliary Parameters	Add any more sshd_config(5) options not covered in this screen. Enter one option per line. These options are case-sensitive. Typos can prevent the SSH service from starting. <small>Close Modal Dialog</small>

Remote systems may require root access. Be sure to have all security precautions in place before allowing *root* access.

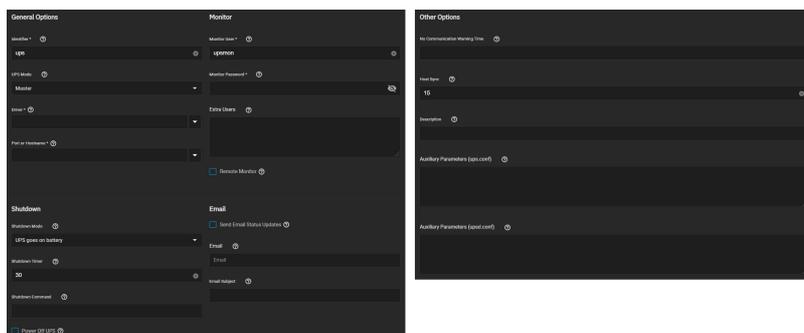
We recommend these additional SSH service options:

- Add `None` to **Auxiliary Parameters** to disable the insecure *none* cipher.
- Increase the **ClientAliveInterval** if SSH connections tend to drop.
- Increase the **ClientMaxStartup** value (**10** is default) when you need more concurrent SSH connections.

Remember to enable the SSH service in **System Settings > Services** after making changes. To create and store specific [SSH connections and keypairs](#), go to **Credentials > Backup Credentials**.

UPS

TrueNAS uses Network UPS Tools [NUT](#) to provide UPS support. After connecting the TrueNAS system UPS device, configure the UPS service by going to **System settings > Services**, finding **UPS**, and clicking



Specific Options	
General Options	
Name	Description
Identifier	Describe the UPS device. It can contain alphanumeric, period, comma, hyphen, and underscore characters.
UPS Mode	Choose Master if the UPS is plugged directly into the system serial port. The UPS will remain the last item to shut down. Choose Slave to have this system shut down before Master. See the Network UPS Tools Overview .
Driver	See the Network UPS Tools compatibility list for a list of supported UPS devices.
Port or Hostname	Serial or USB port connected to the UPS. To automatically detect and manage the USB port settings, select <i>auto</i> . When an SNMP driver is selected, enter the IP address or hostname of the SNMP UPS device.
Monitor	
Name	Description
Monitor User	Enter a user to associate with this service. Keeping the default is recommended.
Monitor Password	Change the default password to improve system security. The new password cannot contain a space or #. Enter accounts that have administrative access. See upsd.users(5) for examples.
Extra Users	Enter accounts that have administrative access. See upsd.users(5) for examples.
Remote Monitor	Set for the default configuration to listen on all interfaces using the known values of user: <code>upsmon</code> and password: <code>fixmepass</code> .
Shutdown	
Name	Description
Shutdown Mode	Choose when the UPS initiates shutdown.
Shutdown Timer	Enter a value in seconds for the the UPS to wait before initiating shutdown. Shutdown will not occur if power is restored while the timer is counting down. This value only applies when Shutdown mode is set to UPS goes on battery.
Shutdown Command	Enter a command to shut down the system when either battery power is low or the shutdown timer ends.
Power off UPS	Set for the UPS to power off after shutting down the system.
Email	
Name	Description
Send Email Status Updates	Set enable sending messages to the address defined in the Email field.
Email	Enter any email addresses to receive status updates. Separate entries by pressing Enter.
Email Subject	Enter the subject for status emails.
Other Options	
Name	Description
No Communication Warning Time	Enter a number of seconds to wait before alerting that the service cannot reach any UPS. Warnings continue until the situation is fixed.
Host Sync	Upsmon will wait up to this many seconds in master mode for the slaves to disconnect during a shutdown situation.
Description	Describe this service.
Auxiliary Parameters (ups.conf)	Enter any extra options from ups.conf .
Auxiliary Parameters (upsd.conf)	Enter any extra options from upsd.conf .

Some UPS models are unresponsive with the default polling frequency (default is **two** seconds). TrueNAS displays the issue in logs as a recurring error like `libusb_get_interrupt: Unknown error`. If you get an error, decrease the polling frequency by adding an entry to **Auxiliary Parameters (ups.conf)**: `pollinterval =`

[upsc\(8\)](#) can get status variables like the current charge and input voltage from the UPS daemon. Run this in **System Settings > Shell** using the syntax `upsc ups@localhost`. The [upsc\(8\)](#) manual page has other usage examples.

[upscmd\(8\)](#) can send commands directly to the UPS, assuming the hardware supports it. Only users with administrative rights can use this command. You can create them in the **Extra Users** field.

How do I find a device name?

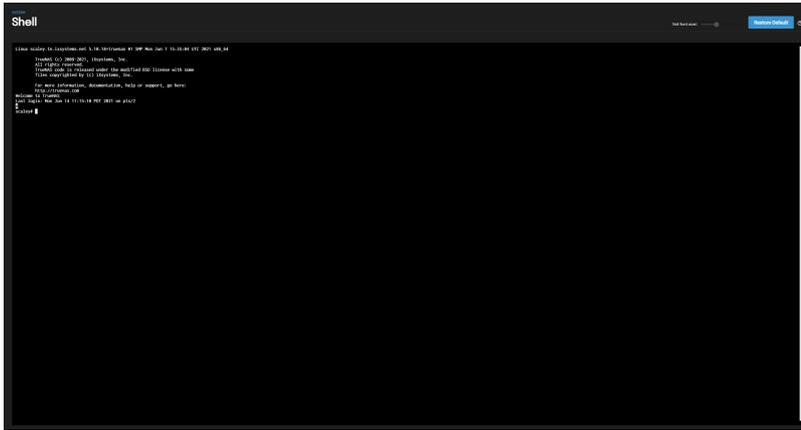
For USB devices, the easiest way to determine the correct device name is to set **Show console messages** in **System Settings > Advanced**. Plug in the USB device and look for a `/dev/ugen` or `/dev/uhid` device name in the console messages.

Can I attach Multiple Computers to One UPS?

A UPS with adequate capacity can power multiple computers. One computer connects to the UPS data port with a serial or USB cable. This primary system makes UPS status available on the network for other computers. The UPS powers the secondary computers, and they receive UPS status data from the primary system. See the [NUT User Manual](#) and [NUT User Manual Pages](#).

12.6 - Shell

The SCALE shell is a convenient means to run command line tools, configure different system settings, or find log files or debug information. The shell opens with the root user logged in.



The *Set font size slider* adjusts the size of text displayed in the shell. Click *RESTORE DEFAULT* to reset the shell font and size.

Shell command history is available for the current session. Use the *Up* and *Down* arrow keys to scroll through previously entered commands. Edit the command if desired, then press *Enter* to re-enter the command. Browsing away from the **Shell** screen clears the command history.

Home, *End*, and *Delete* keys are supported. Tab completion is also available. Type a few letters and press *Tab* to complete a command name or filename in the current directory. Right-clicking in the terminal window displays a reminder about using *Command+c* and *Command+v* or *Ctrl+Insert* and *Shift+Insert* for copy and paste operations in the shell.

Entering *exit* leaves the session. Click *Reconnect* to start a new session.

Default Shell

Clicking other web interface menus closes the shell session and stops commands running in the shell. [zsh](#) is the default shell, but this is changed by editing the *root* user in **Credentials > Local Users** and choosing a different option in the *Shell* drop down. Most Linux command line utilities are available in the shell.

Tmux provides the ability to detach shell sessions and then reattach to them later. Commands continue to run in a detached session.

Experimental CLI

Using the experimental SCALE command line interface (CLI) to directly configure different SCALE features.

This feature is experimental and still in active development. No bug reports or feature requests are being accepted at this time.

To switch to the experimental CLI, enter *cli*. Basic commands are:

- *..* - up one level
- *exit* - exit the CLI
- *ls* - list the available directories and commands
- *? or help* - list the built-in commands

The CLI features an auto-suggest mechanism for commands. Begin typing a command and the CLI shows a list of all matching commands.



The CLI is intended to be an alternative method to configuring TrueNAS features. Because of the variety of available features and configuration, CLI-specific instructions are included in their respective section of the UI documentation.

13 - SCALE API

You can access TrueNAS SCALE API documentation in the web interface by clicking

> **API Keys** > **DOCS**.



Alternatively, append `/api/docs/` to your TrueNAS hostname or IP address in a browser to access the API documentation.

For convenience, we store static builds of the current 2.0 API documentation on the Docs Hub:

- [Websocket Protocol](#)
- [RESTful](#)

14 - Notices

Official statements from iXsystems, Inc.

14.1 - TrueNAS SCALE EULA

TrueNAS SCALE End User License Agreement

Important - Please Read This EULA Carefully

PLEASE CAREFULLY READ THIS END USER LICENSE AGREEMENT (EULA) BEFORE CLICKING THE AGREE BUTTON. THIS AGREEMENT SERVES AS A LEGALLY BINDING DOCUMENT BETWEEN YOU AND IXSYSTEMS, INC. BY CLICKING THE AGREE BUTTON, DOWNLOADING, INSTALLING, OR OTHERWISE USING TRUENAS SCALE SOFTWARE, YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS IN THIS AGREEMENT, DO NOT USE OR INSTALL TRUENAS SCALE SOFTWARE.

This agreement is provided in accordance with the Commercial Arbitration Rules of the American Arbitration Association (the "AAA Rules") under confidential binding arbitration held in Santa Clara County, California. To the fullest extent permitted by applicable law, no arbitration under this EULA will be joined to an arbitration involving any other party subject to this EULA, whether through class arbitration proceedings or otherwise. Any litigation relating to this EULA shall be subject to the jurisdiction of the Federal Courts of the Northern District of California and the state courts of the State of California, with venue lying in Santa Clara County, California. All matters arising out of or relating to this agreement shall be governed by and construed in accordance with the internal laws of the State of California without giving effect to any choice or conflict of law provision or rule.

1.0 Definitions

- 1.1 "Company", "iXsystems" and "iX" means iXsystems, Inc., on behalf of themselves, subsidiaries, and affiliates under common control.
- 1.2 "TrueNAS SCALE Software" means the TrueNAS SCALE storage management software.
- 1.3 "TrueNAS Device" means the TrueNAS storage appliances and peripheral equipment provided by iXsystems or a third party.
- 1.4 "Product" means, individually and collectively, the TrueNAS SCALE Software and the TrueNAS Device provided by iXsystems.
- 1.5 "Open Source Software" means various open source software components licensed under the terms of applicable open source license agreements, each of which has its own copyright and its own applicable license terms.
- 1.6 "Licensee", "You" and "Your" refers to the person, organization, or entity that has agreed to be bound by this EULA including any employees, affiliates, and third party contractors that provide services to You.
- 1.7 "Agreement" refers to this document, the TrueNAS End User License Agreement.

2.0 License

Subject to the terms set forth in this Agreement, iXsystems grants You a non-exclusive, non-transferable, perpetual, limited license without the option to sublicense, to use TrueNAS SCALE Software on Your TrueNAS Device(s). This use includes but is not limited to using or viewing the instructions, specifications, and documentation provided with the Product.

TrueNAS SCALE software is made available as Open Source Software, subject to the license conditions contained within that Open Source Software.

3.0 License Restrictions

TrueNAS SCALE Software is authorized for use on any TrueNAS Device. TrueNAS Devices can include hardware provided by iXsystems or third parties. TrueNAS Devices may also include virtual machines and cloud instances. TrueNAS SCALE software may not be commercially distributed or sold without an addendum license agreement and express written consent from iXsystems.

The TrueNAS SCALE Software is protected by copyright laws and international treaties, as well as other intellectual property laws, statutes, and treaties. The TrueNAS SCALE Software is licensed, not sold to You, the end user. You do not acquire any ownership interest in the TrueNAS SCALE Software, or any other rights to the TrueNAS SCALE Software, other than to use the TrueNAS SCALE Software in accordance with the license granted under this Agreement, subject to all terms, conditions, and restrictions. iXsystems reserves and shall retain its entire right, title, and interest in and to the TrueNAS SCALE Software, and all intellectual property rights arising out of or relating to the TrueNAS SCALE Software, subject to the license expressly granted to You in this Agreement.

The TrueNAS SCALE Software may contain iXsystems' proprietary trademarks and collateral. By agreeing to this license agreement for TrueNAS SCALE, You agree to use reasonable efforts to safeguard iXsystems' intellectual property and hereby agree to not use or distribute iXsystems' proprietary intellectual property and collateral commercially without the express written consent of iXsystems. Official iXsystems Channel Partners are authorized to use and distribute iXsystems' intellectual property through an addendum to this license agreement. By accepting this Agreement, You are responsible and liable for all uses of the Product through access thereto provided by You, directly or indirectly.

The TrueNAS SCALE software includes Open Source components and some proprietary extensions which are available through additional licences You agree to not alter the source code to take advantage of the proprietary extensions without a license to those proprietary extensions, including the TrueNAS Enterprise features sets.

4.0 General

- 4.1 Entire Agreement - This Agreement, together with any associated purchase order, service level agreement, and all other documents and policies referenced herein, constitutes the entire and only agreement between You and iXsystems for use of the TrueNAS SCALE Software and all other prior negotiations, representations, agreements, and understandings are superseded hereby. No agreements altering or supplementing the terms hereof may be made except by means of a written document signed by Your duly authorized representatives and those of iXsystems.
- 4.2 Waiver and Modification - No failure of either party to exercise or enforce any of its rights under this EULA will act as a waiver of those rights. This EULA may only be modified, or any rights under it waived, by a written document executed by the party against which it is asserted.
- 4.3 Severability - If any provision of this EULA is found illegal or unenforceable, it will be enforced to the maximum extent permissible, and the legality and enforceability of the other provisions of this EULA will not be affected.
- 4.4 United States Government End Users - For any TrueNAS SCALE Software licensed directly or indirectly on behalf of a unit or agency of the United States Government, this paragraph applies. Company's proprietary software embodied in the Product: (a) was developed at private expense and is in all respects Company's proprietary information; (b) was not developed with government funds; (c) is Company's trade secret for all purposes of the Freedom of Information Act; (d) is a commercial item and thus, pursuant to Section 12.212 of the Federal Acquisition Regulations (FAR) and DFAR Supplement Section 227.7202, Government's use, duplication or disclosure of such software is subject to the restrictions set forth by the Company and Licensee shall receive only those rights with respect to the Product as are granted to all other end users.
- 4.5 Title - iXsystems retains all rights, titles, and interest in TrueNAS SCALE Software and all related copyrights, trade secrets, patents, trademarks, and any other intellectual and industrial property and proprietary rights, including registrations, applications, registration keys, renewals, and extensions of such rights. Contact Information - If You have any questions about this Agreement, or if You want to contact iXsystems for any reason, please email legal@ixsystems.com.
- 4.6 Maintenance and Support - You may be entitled to support services from iXsystems after purchasing a Product or a support contract. iXsystems will provide these support services based on the length of time of the purchased support contract. This maintenance and support is only valid for the length of time that You have purchased with Your Product. iXsystems may from time to time and at their sole discretion vary the terms and conditions of the maintenance and support agreement based on different business environmental and personnel factors. Any variations will be notified via email and the support portal. For more information on our Maintenance and Support contract, refer to <https://www.ixsystems.com/support/>.
- 4.7 Force Majeure - iXsystems will not be deemed to be in default of any of the provisions of this Agreement or be liable for any delay or failure in performance due to Force Majeure, which shall include without limitation acts of God, earthquake, weather conditions, labor disputes, changes in law, regulation or government policy, riots, war, fire, epidemics, acts or omissions of vendors or suppliers, equipment failures, transportation difficulties, malicious or criminal acts of third parties, or other occurrences which are beyond iXsystems' reasonable control.
- 4.8 Termination - iXsystems may cease any and all support, services, or maintenance under this Agreement without prior notice, or liability, and for any reason whatsoever, without limitation, if any of the terms and conditions of this Agreement are breached. Other provisions of this Agreement will survive termination including, without limitation, ownership provisions, warranty disclaimers, indemnity, and limitations of liability.
- 4.9 Open Source Software Components - iXsystems uses Open Source Software components in the development of the TrueNAS SCALE Software. Open Source Software components that are used in the TrueNAS SCALE Software are composed of separate components each having their own trademarks, copyrights, and license conditions.
- 4.10 Assignment - Licensee shall not assign or otherwise transfer any of its rights, or delegate or otherwise transfer any of its obligations or performance,

under this Agreement, in each case whether voluntarily, involuntarily, by operation of law, or otherwise, without iXsystems' prior written consent. No delegation or other transfer will relieve Licensee of any of its obligations or performance under this Agreement. Any purported assignment, delegation, or transfer in violation of this Section is void. iXsystems may freely assign or otherwise transfer all or any of its rights, or delegate or otherwise transfer all or any of its obligations or performance, under this Agreement without Licensee's consent. This Agreement is binding upon and inures to the benefit of the parties hereto and their respective permitted successors and assigns.

5.0 Export Control Regulations

"The Product may be subject to export control laws. You shall not, directly or indirectly, export, re-export, or release the Product to, or make the Product accessible from, any jurisdiction or country to which export, re-export, or release is prohibited by law, rule, or regulation. You shall comply with all applicable laws, regulations, and rules, and complete all required undertakings (including obtaining any necessary export license or other governmental approval)."

6.0 Data Collection and Privacy

TrueNAS SCALE Software may collect non-sensitive system information relating to Your use of the Product, including information that has been provided directly or indirectly through automated means. Usage of TrueNAS SCALE Software, device status and system configuration are allowed according to iXsystems' privacy policy.

TrueNAS SCALE Software will not collect sensitive User information including email addresses, names of systems, pools, datasets, folders, files, credentials.

By accepting this Agreement and continuing to use the Product, you agree that iXsystems may use any information provided through direct or indirect means in accordance with our privacy policy and as permitted by applicable law, for purposes relating to management, compliance, marketing, support, security, update delivery, and product improvement.

7.0 Limitation of Liability and Disclaimer of Warranty

THE PRODUCT IS PROVIDED "AS IS" AND WITH ALL FAULTS AND DEFECTS WITHOUT WARRANTY OF ANY KIND. TO THE MAXIMUM EXTENT PERMITTED UNDER APPLICABLE LAW, IXSYSTEMS, ON ITS OWN BEHALF AND ON BEHALF OF ITS AFFILIATES AND ITS AND THEIR RESPECTIVE LICENSORS AND SERVICE PROVIDERS, EXPRESSLY DISCLAIMS ALL WARRANTIES, WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE, WITH RESPECT TO THE PRODUCT, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NON-INFRINGEMENT, AND WARRANTIES THAT MAY ARISE OUT OF COURSE OF DEALING, COURSE OF PERFORMANCE, USAGE, OR TRADE PRACTICE. WITHOUT LIMITATION TO THE FOREGOING, IXSYSTEMS PROVIDES NO WARRANTY OR UNDERTAKING, AND MAKES NO REPRESENTATION OF ANY KIND THAT THE PRODUCT WILL MEET THE LICENSEE'S REQUIREMENTS, ACHIEVE ANY INTENDED RESULTS, BE COMPATIBLE, OR WORK WITH ANY OTHER SOFTWARE, APPLICATIONS, SYSTEMS, OR SERVICES, OPERATE WITHOUT INTERRUPTION, MEET ANY PERFORMANCE OR RELIABILITY STANDARDS OR BE ERROR FREE, OR THAT ANY ERRORS OR DEFECTS CAN OR WILL BE CORRECTED.

TO THE FULLEST EXTENT PERMITTED UNDER APPLICABLE LAW: (A) IN NO EVENT WILL IXSYSTEMS OR ITS AFFILIATES, OR ANY OF ITS OR THEIR RESPECTIVE LICENSORS OR SERVICE PROVIDERS, BE LIABLE TO LICENSEE, LICENSEE'S AFFILIATES, OR ANY THIRD PARTY FOR ANY USE, INTERRUPTION, DELAY, OR INABILITY TO USE THE PRODUCT; LOST REVENUES OR PROFITS; DELAYS, INTERRUPTION, OR LOSS OF SERVICES, BUSINESS, OR GOODWILL; LOSS OR CORRUPTION OF DATA; LOSS RESULTING FROM SYSTEM OR SYSTEM SERVICE FAILURE, MALFUNCTION, OR SHUTDOWN; FAILURE TO ACCURATELY TRANSFER, READ, OR TRANSMIT INFORMATION; FAILURE TO UPDATE OR PROVIDE CORRECT INFORMATION; SYSTEM INCOMPATIBILITY OR PROVISION OF INCORRECT COMPATIBILITY INFORMATION; OR BREACHES IN SYSTEM SECURITY; OR FOR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, EXEMPLARY, SPECIAL, OR PUNITIVE DAMAGES, WHETHER ARISING OUT OF OR IN CONNECTION WITH THIS AGREEMENT, BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, REGARDLESS OF WHETHER SUCH DAMAGES WERE FORESEEABLE AND WHETHER OR NOT IXSYSTEMS WAS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES; (B) IN NO EVENT WILL IXSYSTEMS' AND ITS AFFILIATES', INCLUDING ANY OF ITS OR THEIR RESPECTIVE LICENSORS' AND SERVICE PROVIDERS', COLLECTIVE AGGREGATE LIABILITY UNDER OR IN CONNECTION WITH THIS AGREEMENT OR ITS SUBJECT MATTER, UNDER ANY LEGAL OR EQUITABLE THEORY, INCLUDING BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), STRICT LIABILITY, AND OTHERWISE, EXCEED THE TOTAL AMOUNT PAID TO IXSYSTEMS PURSUANT TO THIS AGREEMENT FOR THE PRODUCT THAT IS THE SUBJECT OF THE CLAIM; (C) THE LIMITATIONS SET FORTH IN THIS SECTION SHALL APPLY EVEN IF THE LICENSEE'S REMEDIES UNDER THIS AGREEMENT FAIL OF THEIR ESSENTIAL PURPOSE.

You hereby acknowledge that you have read and understand this Agreement and voluntarily accept the duties and obligations set forth herein by clicking accept on this Agreement.

14.2 - ZFS Feature Flags Removal

Early testers of TrueNAS SCALE are advised:

On June 29, 2021, a new feature was merged into the TrueNAS fork of OpenZFS[1] for developers to test and integrate with other parts of the system. This feature included a new pool feature flag to signify an on-disk format change to how xattr names are encoded on Linux. This original version of the feature was easily activated by a default pool configuration. We quickly decided that the default configuration should not activate this feature until it is available in upstream OpenZFS, and on July 15 we merged changes[2] which make the defaults prevent activation of the new feature.

[1]: <https://github.com/truenas/zfs/pull/8>

[2]: <https://github.com/truenas/zfs/pull/16>

The new feature fixes a long standing issue in ZFS on Linux, which had from its start encoded xattr names in a way that is incompatible with ZFS implementations for every other platform. As one of the planned features of TrueNAS SCALE is the easy migration of pools from TrueNAS CORE, we have been developing this and other missing features to improve feature parity and compatibility across all platforms in OpenZFS. A pull request[3] for the xattr compatibility feature was opened with a request for comments in OpenZFS on April 20, 2021.

[3]: <https://github.com/openzfs/zfs/pull/11919>

On October 6, 2021, we received feedback that the feature flag will not be needed, as a bump to the ZFS POSIX Layer version number should be sufficient. As a result, we have removed the feature flag in question from TrueNAS SCALE to prevent the feature from being enabled moving forward in the release cycle. This is an unfortunate time to receive this insight, as nightly and now beta users of SCALE will have pools created or upgrade with this flag. The impact for most users is negligible, as the pool is still fully operational with the feature flag enabled, as long as it is not active. These users will merely see the unsupported feature is present but inactive:

```
root@truenas[~]# zpool status storage
pool: storage
state: ONLINE
scan: resilvered 860K in 00:00:00 with 0 errors on Mon Oct  4 10:54:52 2021
config:
  NAME                STATE             READ WRITE CKSUM
  storage              ONLINE           0   0   0
  raidz2-0            ONLINE           0   0   0
  99bfb03-cf73-44d4-8418-71b770be1a88  ONLINE           0   0   0
  5c219c07-b29f-45fd-84e8-8e0438dc91c6  ONLINE           0   0   0
  93e7106c-b988-45f4-8437-981b085bcb2  ONLINE           0   0   0
  71de1ea8-5cab-45dd-ae5b-1b19e24ed4b2  ONLINE           0   0   0
  7e75c4d0-c865-40ab-b927-cd13de92fbeb  ONLINE           0   0   0
  c601b128-65e7-4958-884b-235f739d9986  ONLINE           0   0   0
errors: No known data errors
root@truenas[~]# zpool get all storage | grep xattr_compat
storage unsupported@com.ixsystems:xattr_compat inactive      local
root@truenas[~]#
```

Users who created or upgraded a pool using a TrueNAS SCALE build from between June 29 and July 15 2021 or who have manually set `xattr_compat=all` on a dataset and written an xattr will have activated the feature. Once activated, the feature cannot be deactivated until all datasets (including snapshots) that have ever utilized the feature (writing an xattr with `xattr_compat=all` on Linux) have been destroyed. This can be hard to determine, as there is currently no way of checking the feature activation status of a dataset. Most people who did unwittingly activate the feature will merely see the new default value of `xattr_compat=linux` when checking the property.

The feature was marked as read-only compatible, so pools with the feature active are able to be imported read-only on versions of ZFS that do not support the feature. Users are advised to check if their pool has the feature active, and if so, the pool must be backed up and recreated on a version of ZFS without the feature. Builds of SCALE as of October 9, 2021 have the feature removed.

This pool has `feature@xattr_compat` enabled but not active, and can continue to be used on newer versions of TrueNAS SCALE and other ZFS systems:

```
root@truenas[~]# zfs get feature@xattr_compat p0
NAME PROPERTY VALUE SOURCE
p0 feature@xattr_compat enabled local
```

Changing the `xattr_compat` property and writing an xattr in the user namespace activates the feature, preventing the pool from being used on TrueNAS SCALE and other ZFS systems moving forward. The feature is only activated by writing an xattr in the user namespace with `xattr_compat=all` on Linux. Once activated, it stays active even if `xattr_compat=linux` is restored and the file removed:

```
root@truenas[~]# sudo zfs set xattr_compat=all p0
root@truenas[~]# zpool get feature@xattr_compat p0
NAME PROPERTY VALUE SOURCE
p0 feature@xattr_compat enabled local

root@truenas[~]# sudo touch /p0/somefile
root@truenas[~]# zpool get feature@xattr_compat p0
NAME PROPERTY VALUE SOURCE
p0 feature@xattr_compat enabled local

root@truenas[~]# sudo attr -s someattr -V someval /p0/somefile
Attribute "someattr" set to a 7 byte value for /p0/somefile:
someval

root@truenas[~]# zpool get feature@xattr_compat p0
NAME PROPERTY VALUE SOURCE
p0 feature@xattr_compat active local

root@truenas[~]# sudo zfs set xattr_compat=linux p0
root@truenas[~]# sudo rm /p0/somefile
root@truenas[~]# zpool get feature@xattr_compat p0
NAME PROPERTY VALUE SOURCE
p0 feature@xattr_compat active local
```

Creating a new pool with the feature explicitly disabled and replicating the desired datasets is one workaround if your pool has the feature active:

```
root@truenas[~]# sudo zpool create -o feature@xattr_compat=disabled p1 ~/d1
root@truenas[~]# zpool get all p1 | grep xattr_compat
p1 feature@xattr_compat disabled local

root@truenas[~]# sudo zfs snap p0@snap
root@truenas[~]# sudo zfs send p0@snap | sudo zfs recv -F p1
root@truenas[~]# zpool get all p1 | grep xattr_compat
p1 feature@xattr_compat disabled local

root@truenas[~]# zfs get xattr_compat p1
NAME PROPERTY VALUE SOURCE
p1 xattr_compat linux default
```

Please keep in mind these are simplified, contrived examples. If you aren't sure of how to replicate your pool yourself, seek help on the [TrueNAS forums](#).

After upgrade to 22.02-RC.1, the only visible artifact of the feature is that the unsupported flag is present in `zpool get all`:

```
root@truenas[~]# zpool get all storage | grep xattr_compat
storage unsupported@com.ixsystems:xattr_compat inactive      local
```

The unsupported feature will not be presented by `zpool status`.

NOTE: It is not possible to disable the feature once it is enabled; however, having the feature in the enabled state, should not cause a problem. The problem arises when the feature is active. There is currently no practical way to tell which datasets or snapshots are keeping the feature active, so while destroying all traces of it should in theory return the feature from active back to enabled, in practice it's hard to know you won't have to end up destroying the whole pool anyway. For information on how to perform data protection procedures, please refer to the TrueNAS SCALE [Data Protection](#) documentation.

15 - SCALE Security Reports

TrueNAS SCALE is not currently an enterprise release. We only recommended SCALE for early adopters who have a backup plan.

The SCALE 22.02 Security report is available [here](#).

16 - User Recommendations

Because TrueNAS is both Open Source and complicated, the massive user community often creates recommendations for specific hardware or environments. User-created recommendations can be added in this location, but be aware these are provided “as-is” and are not officially supported by iXsystems, Inc.

16.1 - Hardened Backup Repository for Veeam

Abstract

This guide explains in details how to create a **Hardened Backup Repository** for [VeeamBackup](#) with **TrueNAS Scale** that means a repository that will survive to **any remote attack**.

The main idea of this guide is **the disabling of the webUI** with an initialisation script and a cron job **to prevent remote deletion of the ZFS snapshots** that guarantee data immutability.

The key points are:

- Rely on ZFS snapshots to **guarantee data immutability**
- Reduce the surface of attack to the minimum
- When the setup is finished, disable **all remote management interfaces**
- Remote deletion of snapshots is impossible **even if all the credentials are stolen**.
- The only way to delete the snapshot is having **physically access to the TrueNAS Server Console**.

This article targets specifically *TrueNAS Scale* and *Veeam Backup*, but it may also apply to some extent to [TrueNAS Core](#) and/or other backup software.

Installation

Install *TrueNAS Scale 22.02* on a **physical** machine.

- If possible the computer should have at least 2 network interfaces:
 - one dedicated network interface for the management
 - the other one for the data sharing

A virtualized TrueNAS server is not suitable for a hardened backup repository because a malware can easily take the control of TrueNAS server and destroy its data after compromising the hypervisor.

Create a ZFS pool

Go to *Storage* | *Create Pool*

- Name: **tank1**

Even if you can use any pool name, the guide is easier to follow if you use **tank1** as pool name.

- Click on *SUGGEST LAYOUT* to let TrueNAS guessing the best layout for you. In most situations, it will just work very well.
- Review the proposed layout, then click on *CREATE*

For a backup repository, the following layouts will provide a good balance between IOPS, available space and level of redundancy:

- 2 to 4 disks: Stripe of mirrors
- 6 disks: RaidZ2
- 8 to 11 disks: RaidZ3
- 12 disks and more: Stripe of Raidz2/Raidz3

Configure SMART Tests

[SMART](#) (*Self-Monitoring, Analysis and Reporting Technology*) is a monitoring system included in hard disk drives to anticipate imminent hardware failures.

Go to *Data Protection* | *S.M.A.R.T Test* | *Add*

- All Disks
- Type: **LONG**
- Description: **Long SMART test**
- Schedule: **Monthly (0 0 1 * *) on the first day of the month at 00:00 (12:00 AM)**
- SAVE

Configure the network

For a hardened repository, it is better to use a **fixed IP address** than a DHCP configuration, because a compromised DHCP server can provide malicious DNS settings.

Global Network Configuration

Go to *Network* | *Global Configuration*

- *Hostname and Domain*
 - Configure **Hostname** and **Domain**
- *Service Announcement*
 - NetBIOS-NS
 - mDNS
 - WS-Discovery

For a hardened repository it is preferable to disable any service announcement

- *DNS Servers*
 - Nameserver 1: **1.1.1.1**
 - Nameserver 2: **8.8.8.8**

For a hardened server, it is preferable to use the IP addresses of very well known and secure public DNS than your own internal DNS server.

- Cloudflare: 1.1.1.1
- Google: 8.8.8.8

- *Default Gateway*
 - Setup *IPv4 (or IPv6) Default Gateway* according to your network
- *Outbound Network*
 - **(o) Allow Specific**
 - Enable **Mail** and **Update**
- Other Settings
 - *HTTP Proxy*: stay empty

Connecting to Internet through a proxy is a good security practice because it prevents malwares to communicate easily with their control and command servers, but it is out of the scope of this guide.

- SAVE

Network Interfaces Configuration

Go to *Network | Interfaces*

- Click on the first interface and configure it as the management interface

- Management interface

- **Description: management**
- DHCP
- Autoconfigure IPv6
- **Other Settings**
 - Disable Hardware Offloading
 - **MTU: 1500**

For a hardened repository, it is preferable to keep the default value (1500) for the MTU, because using jumbo frame makes the network configuration more complex to manage.

- **IP Addresses**
 - Add the IP address of the management interface
- **APPLY**
- **TEST CHANGES**

When you are testing the new network settings, you have 60 seconds to confirm that it works by clicking on **SAVE CHANGES**, otherwise the system automatically rolls back to the previous network configuration to avoid kicking you out of the network.

- Data interface

- Management interface
 - **Description: data sharing**
 - DHCP
 - Autoconfigure IPv6
 - **Other Settings**
 - Disable Hardware Offloading
 - **MTU: 1500**
 - **IP Addresses**
 - Add the IP address of the data sharing interface
- **APPLY**
- **TEST CHANGES**
- **SAVE CHANGES**

Configure the user accounts

Setup root account

Go to *Credentials | Local Users*

- Edit the *root* user
 - Fill the *Email* field

System notification are sent by email to the **root** user, so this email address is very important.

- If you wish to use SSH for management, fill also *SSH Public Key*

SSH is more convenient than the web shell interface to enter commands that are missing from the web user interface.

Create a account for Veeam

Go to *Credentials | Local Groups | Add*

- **GID: 10000**
- **Name: veeam**
- Permit Sudo
- Samba Authentication
- Allow Duplicated GIDs
- SAVE

Go to *Credentials | Local Users | Add*

- **Full Name: Veeam Backup**
- **Username: veeam**
- **Password: use a very long and strong password**
- **Password confirmation:**
- **Email:** stay empty
- **User ID and Groups**
 - **User ID: 10000**
 - New Primary Group
 - **Primary Group: veeam**
 - **Auxiliary group:** stay empty
- **Directories and Permissions**
 - **Home Directory: /nonexistent**
 - **Home Directory Permission: clear all permissions, except user permissions**
 - **SSH Public Key:** stay empty
 - **Disable password: no**
 - **Shell: nologin**
 - Lock User
 - Permit Sudo
 - Microsoft Account
 - Samba Authentication
- SAVE

Configure SSH

Go to *System Settings | Services | SSH* and click on the pencil (✎)

- Click **ADVANCED SETTINGS**
 - **TCP Port: 22**
 - Log in As Root with Password
 - Allow Password Authentication
 - Allow Kerberos Authentication
 - Allow TCP Port Forwarding
 - **Bind Interfaces: use the management network interface**
 - Compress Connections
 - SFTP Log Level: stay blank
 - SFTP Log Facility: stay blank
 - Weak Ciphers: **None, AES128-CBC**
 - Auxiliary Parameters: **AllowUsers root@192.168.0.10**
 - where 192.168.0.10 is the IP address of your desktop computer you use to manage the TrueNAS server.

- **SAVE**
- Toggle the running button to start the SSH service **but do not start automatically SSH**

Do not start automatically SSH because we will disable the SSH service later to harden the repository.

Configure the mail notification

Configuring the mail notification is very important, because it will be the only way to know that happens (for example if a disk is dying) after disabling the web management interface to harden the repository.

Edit mail notification

- Click on the **bell** 🔔 icon on the top right corner
- Click on the **gear** ⚙ icon
- Select **Email**
- Fill the web form according to your email provider
- **Send Test Mail**
- Check that you receive the testing email
- **SAVE**

Create a dataset for Veeam

Go to *System Settings* | *Shell* (or connect with SSH)

```
zfs create tank1/veeam
zfs set org.freenas:description="veeam hardened repo" tank1/veeam
zfs set compression=off tank1/veeam
chown veeam:veeam /mnt/tank1/veeam
chmod 700 /mnt/tank1/veeam
```

Description of shell commands

1. Create a dataset name **tank1/veeam**
2. Set dataset description ("veeam hardened repo")
3. Set compression level to **off** because Veeam backup are already compressed
4. Set ownership of user **veeam** and group **veeam** on directory **/mnt/tank1/veeam**
5. Set restrictive user permissions on **/mnt/tank1/veeam**

If you really following this guide from scratch, then the dataset **tank1/veeam** is empty, then you can create an **empty snapshot** and **lock it** to prevent deleting by mistake the dataset from the web user interface or with the command `zfs destroy`

```
zfs snap tank1/veeam@LOCKED
zfs hold LOCKED tank1/veeam@LOCKED
```

Description of shell commands

1. Create a snapshot named **LOCKED** on **tank1/veeam**.
2. Hold a lock named **LOCKED** on the snapshot. Indeed The name of the snapshot and the name of the lock can be different, but it is easier to use twice the same name.

More information about ZFS locked snapshot

- To lock a snapshot use `zfs hold LOCK_NAME SNAPSHOT_NAME`
- Snapshot can have multiple locks, each lock must have a different name
- A locked snapshot cannot be deleted
- To unlock a snapshot, use `zfs release LOCK_NAME SNAPSHOT_NAME`
- To list the lock names of a particular snapshot, use `zfs holds SNAPSHOT_NAME`
- A dataset with a locked snapshot cannot be deleted neither with the webui nor with the `zfs destroy` command, so it avoid human errors.

Configure ZFS periodic snapshots

Create 3 periodic (hourly, daily and weekly) ZFS snapshots to recover the data if they are deleted or modified.

Hourly snapshots

Go to **Data Protection** | **Periodic Snapshot Tasks**

- **Dataset** **tank1**
- **Exclude:** stay empty
- **Recursive**
- **Snapshot lifetime:** **1 day**
- **Naming Schema:** **auto-%Y%m%d_%H%M-hourly**
- **Schedule:** **Hourly (0 * * * *) at the start of each hour**
- **Begin:** **00:00:00**
- **End:** **23:59:00**
- Allow Taking Empty Snapshots
- **Enabled**
- **SAVE**

It is easier to setup the periodic snapshot at the root dataset and to enable *recursive* snapshot.

Daily snapshots

Go to **Data Protection** | **Periodic Snapshot Tasks**

- **Dataset** **tank1**
- **Exclude:** stay empty
- **Recursive**
- **Snapshot lifetime:** **1 week**
- **Naming Schema:** **auto-%Y%m%d_%H%M-daily**
- **Schedule:** **Daily (0 0 * * *) at 00:00 (12:00 AM)**
- Allow Taking Empty Snapshots
- **Enabled**
- **SAVE**

Weekly snapshots

Go to **Data Protection** | **Periodic Snapshot Tasks**

- **Dataset** **tank1**
- **Exclude:** stay empty
- **Recursive**
- **Snapshot lifetime:** **1 month**
- **Naming Schema:** **auto-%Y%m%d_%H%M-weekly**

- Schedule: **Weekly (0 0 * * sun) on Sundays at 00:00 (12:00 AM)**
- Allow Taking Empty Snapshots
- Enabled
- SAVE

If you have enough disk space, you can use longer retention time. The longer the snapshot are kept, the better your safety is.

Configure Samba Service

Go to *System Settings* | *Services* | *SMB* and click on the pencil (✎)

- Click **ADVANCED SETTINGS**
 - NetBIOS Name: **strongbox** (you can use any name here)
 - NetBIOS Alias: stay empty
 - Workgroup: **WORKGROUP**
 - Description: **Hardened TrueNAS**
 - Enable SMB1 support
 - NTLMv1 Auth
 - UNIX Charset: **UTF-8**
 - Log Level: **Minimum**
 - Use Syslog Only
 - Local Master
 - Enable Apple SMB2/3 Protocol Extensions
 - Administrators Group: stay empty
 - Guest Account: **nobody**
 - File Mask: **0600**
 - Directory Mask: **0700**
 - Bind IP Address: bind on the IP address of the data network interface
 - Auxiliary Parameters: stay empty
 - SAVE
- Toggle the running button to start the SMB service
- Start Automatically SMB

Configure Samba share for Veeam

Go to *Shares* | *Windows (SMB) Shares* | *ADD*

- Click on ***ADVANCED OPTIONS**
 - Basic
 - Path: **/mnt/tank1/veeam**
 - Name: veeam
 - Purpose: **Multi-protocol (NFSv3/SMB) shares**
 - Description: **hardened veeam repository**
 - Enabled
 - Access
 - Enable ACL
 - Export Read Only
 - Browseable to Network client
 - Allow guest access
 - Allow based shared enumeration
 - Host Allow: *put the IP of the Veeam Software server here*
 - Host Deny: stay empty
 - Other Options
 - Use as home share
 - Timemachine
 - Legacy AFP compatibility
 - Enable shadow copy
 - Export Recycle bin
 - Use Apple-Style Character Encoding
 - Enable alternate data streams
 - Enable SMB2/3 Durable handles
 - Enable FSRVP
 - Path suffix: **stay empty, very important**
 - Auxiliary parameters: stay empty
 - SAVE

Add this repository to Veeam Software

See the [documentation of Veeam Backup](#) to add this SAMBA share as a backup target.

In the Veeam wizard select

- **Network attached storage**
- **SMB Share**
- For the credentials, use the veeam account creates on the hardened backup repository (see above)

Hardened the repository

To hardened the backup repository, just remove any possibility to remotely destroy the ZFS snapshots.

Enable password for console access

Go to *System Settings* | *Advanced* | *Console* | *Configure*

- Show Text Console without Password Prompt
- SAVE

Disconnect IPMI

If your server has a [IPMI](#) interface, **physically disconnect the network cable**.

- If a malware takes the control of your management computer, it can use the IPMI interface to destroy your backups.
- Be cautious and just disconnect the cable.

Check that NTP works as expected

- Go to *System Settings* | *General* | *NTP Servers*

By default TrueNAS Scale comes with the following NTP servers

- 0.debian.pool.ntp.org
- 1.debian.pool.ntp.org
- 2.debian.pool.ntp.org

Open a shell

- Go to *System Settings* | *Shell*
- Enter the command `ntpq -p`
- The output will look like

#	ntpq -p	remote	refid	st	t	when	poll	reach	delay	offset	jitter
*	ntppub.darksky.	172.18.1.20		2	u	326	1024	377	11.447	+0.475	0.531
+	ip139.ip-5-196-	145.238.203.14		2	u	208	1024	377	11.484	-0.249	0.279
+	ns2.euskill.com	193.107.56.120		4	u	33	1024	377	22.541	+0.167	0.538

Do not worry if you have different remote hostnames or IP addresses for NTP servers, it is normal because domain names of **ntp.org** point to a pool of servers.

Configure HTTPS

Create an Internal Certificat Authority

Go to [Credentials](#) | [Certificates](#) | [Certificates Authorities](#) | [Add](#)

- **Identifier and Type**
 - Name: **hardened-truenas-scale-ca**
 - Type: **Internal CA**
 - Profiles: **CA**
- **Certificate Options**
 - Key Type: **RSA**
 - Key Length: **4096**
 - Digest Algorithm: **SHA512**
 - Lifetime: **3650**
- **Certificate Subject**
 - Country: **United States**
 - State: **California**
 - Locality: **San Francisco**
 - Organization: **The Name of My Company**
 - Organization Unit: **Backup Department**
 - Email: firstname.surname@the-name-of-my-company.com
 - Common Name: **hardened-truenas-scale-ca**
 - Subject Alternate Names: **hardened-truenas-scale-ca**
- **Extra Constraints**
 - Basic Constraints
 - Authority Key Identifier
 - Extended Key Usage
 - Key Usage
- **Confirm Options**
 - **SAVE**

Create a certificate for HTTPS

Go to [Credentials](#) | [Certificates](#) | [Certificates](#) | [Add](#)

- **Identifier and Type**
 - Name: **hardened-truenas-scale-cert**
 - Type: **Internal Certificate**
 - Profiles: **---**
- **Certificate Options**
 - Signing Certificate Authority: **hardened-truenas-scale-ca**
 - Key Type: **RSA**
 - Key Length: **4096**
 - Digest Algorithm: **SHA512**
 - Lifetime: **3650** (10 years)
- **Certificate Subject**
 - Country: **United States**
 - State: **California**
 - Locality: **San Francisco**
 - Organization: **The Name of My Company**
 - Organization Unit: **Backup Department**
 - Email: firstname.surname@the-name-of-my-company.com
 - Common Name: **hardened.mydomainname.com** (the full qualified domain name)
 - Subject Alternate Names: **hardened.mydomainname.com** (the full qualified domain name)
- **Extra Constraints**
 - Basic Constraints
 - Authority Key Identifier
 - Extended Key Usage
 - Key Usage
- **Confirm Options**
 - **SAVE**

Apply the new HTTPS certificate

Go to [System Settings](#) | [General](#) | [GUI](#) | [Settings](#)

- **GUI**
 - GUI SSL Certificate : **hardened-truenas-scale-cert**
 - Web Interface IPv4 Address: **select the management interface**
 - Web Interface IPv6 Address: **::**
 - Web Interface HTTP Port: **80**
 - Web Interface HTTPS Port: **443**
 - HTTPS Protocols: **TLSv1.3**
 - Web Interface HTTP -> HTTPS Redirect
- **Other Options**
 - Crash reporting
 - Usage collection
 - Show Console Messages
- **SAVE**

Restart Web Service: **CONFIRM, CONTINUE**

Enable Two-Factor Authentication (2FA)

Two-Factor Authentication is time-based, and requires that the system time is set correctly, so check before that NTP works.

- Install an application on your smartphone to generate an [One-Time-Password](#) from a QR-Code. For example [FreeOTP Authenticator](#)
- Go to [Credentials](#) | [2FA](#)
- Keep the default
 - One-Time Password (OTP) Digits: 6
 - Interval: 30
 - Window: 0
 - Enable Two-Factor Auth for SSH
- Click on **Enable Two-Factor Authentication**
- **SHOW QR**
- Use [FreeOTP](#) to capture the QR code
- Log out the web interface
- Test Two-Factor Authentication
- If something goes wrong you can disable the 2FA from the console

```
midclt call auth.twofactor.update '{"enabled": false}'
```

Disable SSH for normal operations

Letting SSH service running is dangerous: if someone steals your SSH private key and passphrase, he can remotely connect to the backup repository and destroy the data.

Check SSH does not automatically start

Go to [System Settings](#) | [Services](#)

- Check that SSH does not start automatically

Stop SSH service on boot

Add a startup script to stop the SSH service in case it has been enabled by mistake

Go to [System Settings](#) | [Advanced](#) | [Init/Shutdown Scripts](#) | [Add](#)

- **Description:** Stop SSH at startup
- **Type:** Command
- **Command:** `/usr/bin/systemctl stop ssh`
- **When:** Post Init
- Enabled
- **Timeout:** 10
- SAVE

Stop SSH service at midnight

To avoid the SSH service stays enabled forever, stop it automatically at midnight

Go to [System Settings](#) | [Advanced](#) | [Cron Job](#) | [Add](#)

- **Description:** stop ssh at midnight
- **Command:** `/usr/bin/systemctl stop ssh`
- **Run as user:** root
- **Schedule:** `*daily (0 0 * *) at 00:00 (12:AM)`
- hide standard output
- hide standard error
- Enabled
- SAVE

Disable Web User Interface for normal operations

Stop WebUI on boot

Go to [System Settings](#) | [Advanced](#) | [Init/Shutdown Scripts](#) | [Add](#)

- **Description:** Stop webUI at startup
- **Type:** Command
- **Command:** `/usr/bin/systemctl stop nginx`
- **When:** Post Init
- Enabled
- **Timeout:** 10
- SAVE

Stop WebUI at midnight

To avoid the WebUI stays enabled forever, stop it automatically at midnight

Go to [System Settings](#) | [Advanced](#) | [Cron Job](#) | [Add](#)

- **Description:** stop webUIh at midnight
- **Command:** `/usr/bin/systemctl stop nginx`
- **Run as user:** root
- **Schedule:** `*daily (0 0 * *) at 00:00 (12:AM)`
- hide standard output
- hide standard error
- Enabled
- SAVE

Change the message of the day

Go to [System Settings](#) | [Advanced](#) | [Console](#) | [Configure](#)

- MOTD Banner: **Hardened repository without remote management, to enable temporary the web interface type "systemctl start nginx"**
- SAVE

Backup the server configuration

Go to [System Settings](#) | [General](#) | [Manage Configuration](#)

- [DOWNLOAD FILE](#)

Test the setup

Reboot the server to check that the web interface is disabled when the computer boots

Daily management

You can temporary enable the web interface to change the configuration

Enable the web interface

Connect to the console and type:

```
systemctl start nginx
```

If you forgot to stop the webUI when you have finished your work, the cron job will do if for you at midnight

Disable the web interface

To immediately disable the web interface connect to the console and type:

```
systemctl stop nginx
```

Recover data after an attack

If your Veeam backup files have been altered it means that the password to access the SAMBA share has been compromised, so you have to change it immediately.

Change the password for the veeam account

Go to [Credentials](#) | [Local Users](#) | [veeam](#)

- Unroll the options, click **EDIT**
- Change *Password*
- **SAVE**

Lock the snapshot to preserve the data

It may take few day to audit your system after an attack, therefore it is a good idea to lock all snapshots to avoid they are automatically deleted when they reached their end of life.

Run the following command in the shell

```
for s in `zfs list -r -t snap -H -o name tank1/veeam`; do zfs hold LOCKED $s ; done
```

Clone the healthy snapshot

Go to *Storage | Snapshots*

- Pick the healthy snapshot
- Unroll the option
- Click *CLONE TO NEW DATASET*
 - *Name:* **tank1/veeam-snap-clone**
 - **SAVE**

Create a new Samba Share to export the cloned dataset

- Use the above instruction to share **tank1/veeam-snap-clone** with SAMBA.
- Reinstall Veeam on a new server
- Connect to the new SAMBA share
- Restore your data.

- The guide for a hardened repository is finished
- Enjoy your hardened repository, and sleep more peacefully at night.