

TECH NOTE

Nutanix Files

Copyright

Copyright 2021 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

Contents

1. Executive Summary.....	5
2. Audience and Purpose.....	6
3. Nutanix Enterprise Cloud Overview.....	8
Nutanix HCI Architecture.....	9
4. Nutanix Files Architecture.....	11
File Server Virtual Machine.....	11
Exports and Shares.....	16
Load Balancing and Scaling.....	25
High Availability.....	28
Active Directory and SMB Operations.....	31
Active Directory, LDAP, and NFS Operations.....	38
Multiprotocol.....	42
Quotas.....	45
Selective File Blocking.....	47
File Analytics.....	49
Smart Tier.....	57
Hypervisor-Specific Support.....	63
5. Backup and Disaster Recovery.....	65
Self-Service Restore.....	65
Protection Domains and Consistency Groups.....	66
Cluster Migration, Failure, and Restoration.....	67
Cloning.....	67
Files Smart Disaster Recovery.....	68
6. Third-Party Integration.....	74
Antivirus.....	74
File Operations Monitoring.....	77
Intelligent Backup.....	79

7. Conclusion.....	84
Appendix.....	85
About Nutanix.....	85
List of Figures.....	86
List of Tables.....	89

1. Executive Summary

Nutanix Files is a software-defined, scale-out file storage solution that provides a repository for unstructured data, such as home directories, user profiles, departmental shares, application logs, backups, and archives. Flexible and responsive to workload requirements, Files is a fully integrated, core component of Nutanix.

You can deploy Nutanix Files on an existing cluster or a standalone cluster. Unlike standalone NAS appliances, Files consolidates VM and file storage, eliminating the need to create an infrastructure silo. Administrators can manage Files with Nutanix Prism, just like VM services, which unifies and simplifies management. Integration with Active Directory enables support for quotas and access-based enumeration, as well as self-service restores with the Windows Previous Versions feature. Nutanix Files also supports native remote replication and file server cloning, which lets you back up Files off-site and run antivirus scans and machine learning without affecting production.

Nutanix Files can run on a dedicated cluster or be collocated on a cluster running user VMs. Nutanix supports Files with both ESXi and AHV. Files includes native high availability and uses Nutanix distributed storage for intracluster data resiliency. AOS distributed storage also provides data efficiency techniques such as erasure coding (EC-X).

Nutanix Files includes File Analytics, which gives you a variety of useful insights into your data, including full audit trails, anomaly detection, ransomware detection and intelligence, data age analytics, and custom reporting.

2. Audience and Purpose

This tech note is part of the Nutanix Solutions Library. We wrote it for architects and systems engineers who want to use Nutanix Files as a simple way to deliver user and group file management. This document describes how to implement and operate Files in your datacenter.

We cover the following subject areas:

- Overview of the Nutanix architecture with Files.
- Load balancing of standard and distributed shares (SMB) and exports (NFS).
- High availability.
- Backup and recovery.
- Quotas and permission management.
- Antivirus.

Table 1: Document Version History

Version Number	Published	Notes
1.0	December 2016	Original publication.
1.1	May 2017	Updated Backup and Disaster Recovery section.
1.2	September 2017	Updated for version 2.2 features.
2.0	February 2018	Updated for version 3.0.
2.1	April 2018	Solution overview update.
2.2	August 2018	SMB share and NFS export updates.
3.0	October 2018	Updated for version 3.1 and updated product naming.
4.0	January 2019	Updated for version 3.2 and AOS 5.9 and later.

Version Number	Published	Notes
5.0	April 2019	Updated for version 3.5 and AOS 5.10 and later.
5.1	August 2019	Updated for version 3.5.1.
5.2	October 2019	Updated for version 3.6.
5.3	October 2020	Updated for version 3.7
6.0	April 2021	Updated for Files 3.8, File Analytics 3.0, and Files Manager 2.0.
7.0	October 2021	Updated for Files 4.0 and Data Lens. Added Smart Tier and SSR Interoperability sections. Updated Shares and Exports, Notifications, Selective File Block, and Files Operations Monitoring sections. Updated Files Smart DR figure.

3. Nutanix Enterprise Cloud Overview

Nutanix delivers a web-scale, hyperconverged infrastructure solution purpose-built for virtualization and both containerized and private cloud environments. This solution brings the scale, [resilience](#), and economic benefits of web-scale architecture to the enterprise through the Nutanix enterprise cloud platform, which combines the core HCI product families—Nutanix AOS and Nutanix Prism management—along with other software products that automate, secure, and back up cost-optimized infrastructure.

Available attributes of the Nutanix enterprise cloud OS stack include:

- Optimized for storage and compute resources.
- Machine learning to plan for and adapt to changing conditions automatically.
- Intrinsic security features and functions for data protection and cyberthreat defense.
- Self-healing to tolerate and adjust to component failures.
- API-based automation and rich analytics.
- Simplified one-click upgrades and software life cycle management.
- Native file services for user and application data.
- Native backup and disaster recovery solutions.
- Powerful and feature-rich virtualization.
- Flexible virtual networking for visualization, automation, and security.
- Cloud automation and life cycle management.

Nutanix provides services and can be broken down into three main components: an HCI-based distributed storage fabric, management and operational intelligence from Prism, and AHV virtualization. Nutanix Prism furnishes one-click infrastructure management for virtual environments running on AOS. AOS is hypervisor agnostic, supporting two third-party hypervisors

—VMware ESXi and Microsoft Hyper-V—in addition to the native Nutanix hypervisor, AHV.

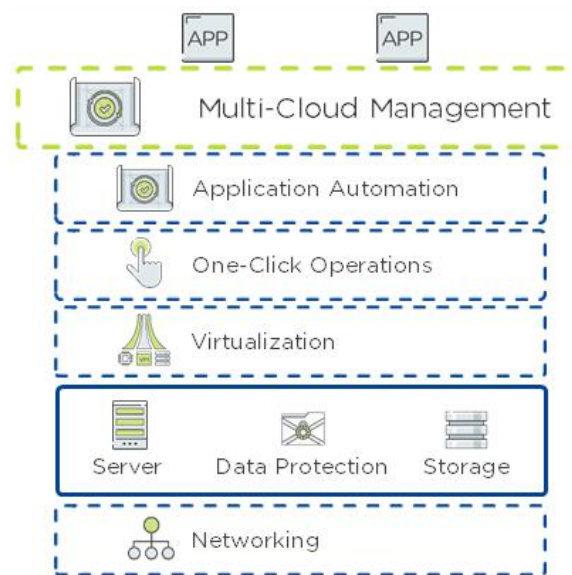


Figure 1: Nutanix Enterprise Cloud OS Stack

Nutanix HCI Architecture

Nutanix does not rely on traditional SAN or network-attached storage (NAS) or expensive storage network interconnects. It combines highly dense storage and server compute (CPU and RAM) into a single platform building block. Each building block delivers a unified, scale-out, shared-nothing architecture with no single points of failure.

The Nutanix solution requires no SAN constructs, such as LUNs, RAID groups, or expensive storage switches. All storage management is VM-centric, and I/O is optimized at the VM virtual disk level. The software solution runs on nodes from a variety of manufacturers that are either entirely solid-state storage with NVMe for optimal performance or a hybrid combination of SSD and HDD storage that provides a combination of performance and additional capacity. The storage fabric automatically tiers data across the cluster to different classes of storage devices using intelligent data placement algorithms. For best

performance, algorithms make sure the most frequently used data is available in memory or in flash on the node local to the VM.

To learn more about Nutanix enterprise cloud software, visit [the Nutanix Bible](#) and [Nutanix.com](#).

4. Nutanix Files Architecture

Nutanix Files is a scale-out approach that provides Server Message Block (SMB) and Network File System (NFS) file services to clients. Nutanix Files instances are composed of a set of VMs (called FSVMs). Files requires at least three FSVMs running on three nodes to satisfy a quorum for high availability.

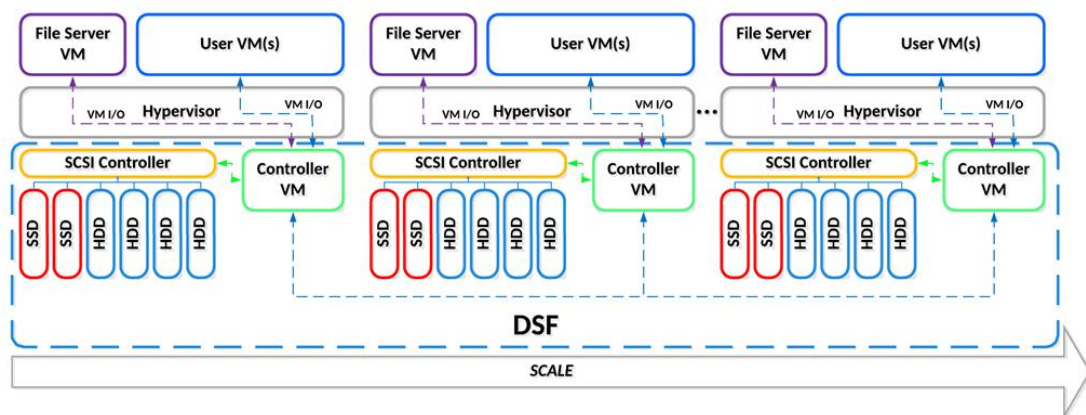


Figure 2: Nutanix Files Instances Run as VMs

File Server Virtual Machine

The File Server VM (FSVM) is based on CentOS and incorporates all the security and hardening that goes into the Nutanix Controller VM (CVM). All the FSVMs have the same basic configuration: four vCPU and 12 GiB of RAM. You can add more vCPU, RAM, and FSVMs to the cluster. For each file server the number of FSVMs must be less than or equal to the number of nodes in the Nutanix cluster; however, you can create multiple file server deployments if needed. Each Nutanix Files cluster can support up to 16 FSVMs. Nutanix Files 3.1 introduced single-FSVM deployments intended for one- and two-node Nutanix clusters. You can also have single-FSVM deployments for larger clusters with AOS 5.10.1 or later.

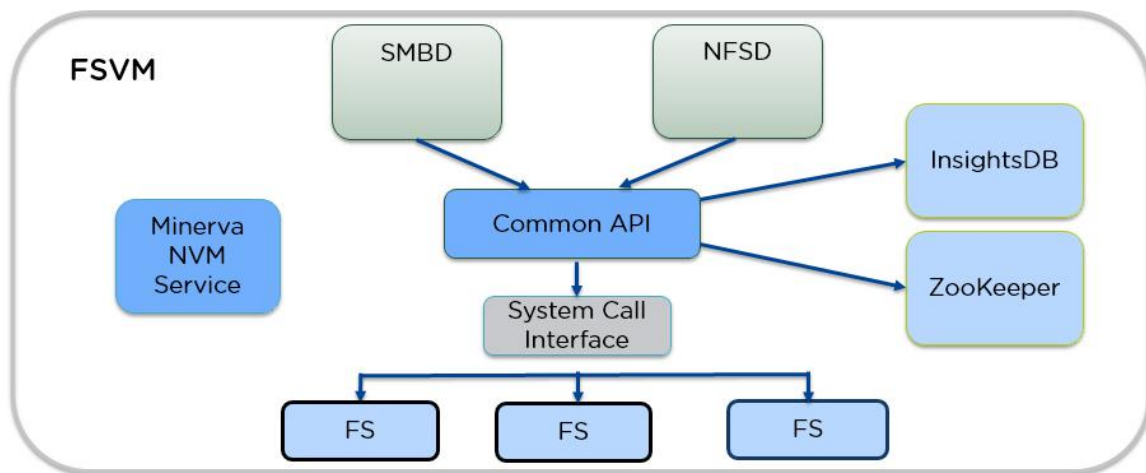


Figure 3: Data Path Architecture of Nutanix Files

Nutanix Files can support SMB and NFS from the same FSVM, and beginning with Files version 3.5 you can enable simultaneous SMB and NFS access to the same share and data, commonly referred to as multiprotocol support. Both SMB and NFS share a common library, allowing a modular approach. InsightsDB is a NoSQL database that maintains the statistics and alerts for Files. Zookeeper is a centralized service that maintains configuration information, such as domain, share or export, and IP information. The Minerva NVM Service talks to the local CVM and sends heartbeats to share health information and help with failover.

Each FSVM stores file server data on multiple file systems that store share-specific data. The individual file system provides the snapshot capability used to provide Windows Previous Versions (WPV) support to clients. By using separate file systems for each share or export, Nutanix Files can scale to support billions of files in one cluster.

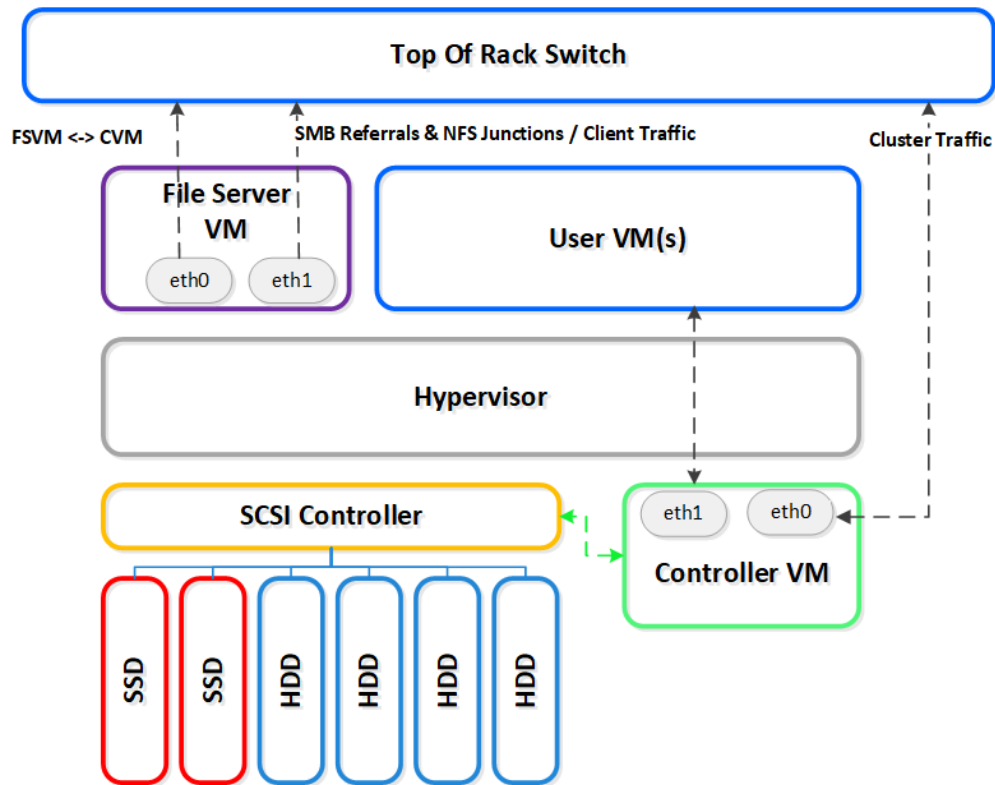


Figure 4: FSVM Internal Communication on One Node

The previous diagram shows one FSVM running on a node, but you can put multiple FSVMs on a node for multitenancy.

Networking

The FSVM has two network interfaces: the storage interface and the client interface. The FSVM service that talks to the CVMs uses the storage interface, which also provides access to Nutanix Volumes iSCSI vDisks in volume groups. The storage interface helps manage deployment, failover, and maintenance and enables control over one-click upgrades. Integration with the CVM lets the FSVM determine if a storage fault has occurred and, if so, whether you must take action. The FSVM service sends a heartbeat to its local CVM service each second indicating its state.

Tip: We recommend that you place the FSVM storage interface on the same network VLAN as the Nutanix CVM management interface (eth0 of the CVM) to help ensure maximum performance for the iSCSI sessions.

The client interface allows clients to connect to SMB shares and NFS exports hosted on an FSVM. A client can connect to any FSVM client network interface to access their file data. If a different FSVM provides the data, the client connection automatically redirects to the correct FSVM interface. If an FSVM fails, the client network address for the failed FSVM moves to another to preserve data access.

Storage

Nutanix File Server VM

As shown in the following figure, each FSVM uses three separate vDisks: a 12 GiB boot disk that contains the boot image, a 45 GiB disk (/home/nutanix) that contains logs and software state, and a 45 GiB disk for Cassandra data.




Figure 5: FSVM vDisks and Volume Groups

Starting with the Files 3.7 release, each FSVM also has a volume group that helps maintain auditing events and persistent handles for SMB Transparent Failover.

Shares and Exports

A Nutanix Files cluster is a single namespace that includes a collection of file systems used to support the SMB and NFS shares. The file systems support dynamic metadata, which enables you to store an unlimited number of files. The file system also supports variable block length allocations up to the default size of 64 KB. The variable block length matches the size of the file; for example, a 1 KB file allocates 1 KB on the file system. Any file over 64 KB allocates in 64

KB increments. Starting with Files 3.7, you can specify the maximum allocation size to improve performance. For environments with more random access patterns, you can choose a maximum allocation size of 16 KB. For environments with sequential patterns, you can choose 1 MB. These file systems use Nutanix volume groups as storage.

Tip: For guidance on whether you should use the random or sequential file system option, see the [Nutanix Files Performance Tech Note](#) (requires Nutanix Portal account).

Volume groups enable both high availability and scale-out for the file systems. A volume group is a collection of logically related vDisks (or volumes) attached to the guest via iSCSI. When an FSVM is down for maintenance or a fault occurs, one of the surviving FSVMs takes over volume group ownership and continues servicing requests.

With the 3.x release of Nutanix Files, volume groups contain 16 vDisks. Files 3.x versions prior to 3.2 use 2 disks for metadata and 4 disks for data, as in the 2.x release, with the remaining disks reserved for expansion in later versions. From the 3.2 release of Files onward, all disks in the volume group can be used by the assigned shares.

With the 4.x release of Nutanix Files, volume groups contain 17 vDisks. We added an additional vDisk as a separate intent log (SLOG) to accept synchronous writes, as the SLOG device helps with long-duration synchronous random write performance.

The initial file system pool created on the volume group uses six of the drives noted and supports 40 TB of data.



Figure 6: Initial 40 TB File System Pool

Once this file system pool reaches 80 percent space utilization, Nutanix Files automatically expands it to 80 TB by incorporating another four data disks from the volume group.



Figure 7: 80 TB File System Pool Following First Expansion

This expansion process can occur two more times, increasing the file system capacity to 120 TB with the second expansion and 140 TB with the final expansion. Each expansion triggers when you've used 80 percent of the space in the current pool.




Figure 8: 120 TB File System Pool After Second Expansion



Figure 9: 140 TB File System Pool After Final Expansion

The Nutanix Files 3.7 release introduced a scale-up function to the file system pool. Once you've scaled a file system out to use all 16 disks, if that pool reaches 80 percent capacity, it uses additional space on the existing disks in the volume group. This vertical expansion allows support for a file system pool of up to 280 TB.

In summary, the maximum sizes for individual volume groups based on Nutanix Files version are as follows:

- Nutanix Files 2.x to 3.2: 40 TB
- Nutanix Files 3.2 to 3.6: 140 TB
- Nutanix Files 3.7 forward: 280 TB

Tip: You must have created the File server with a minimum version of 3.x to take advantage of horizontal scale-out to 140 TB and vertical scale-up to 280 TB.

Exports and Shares

There are two types of SMB shares and NFS exports.

SMB shares:

1. Distributed (previously called home).
2. Standard (previously called general).

NFS exports:

1. Distributed (previously called sharded).
2. Standard (previously called nonsharded).

A standard share is an SMB share or NFS export hosted by a single FSVM. A distributed share spreads the workload by distributing the hosting of top-level directories across FSVMs, which also simplifies administration.

Starting with Files version 3.7, you can store files in the root of distributed NFS shares. Distributed SMB shares don't allow files in the root.

Distributed Shares

Distributed shares and exports distribute data by dividing the top-level directories across all the FSVMs that make up the file server. Nutanix Files maintains the directory mapping for each responsible FSVM using InsightsDB. FSVMs use distributed file system (DFS) referrals for SMB and junctions for NFSv4 to make sure that clients can connect to the right top-level directories.

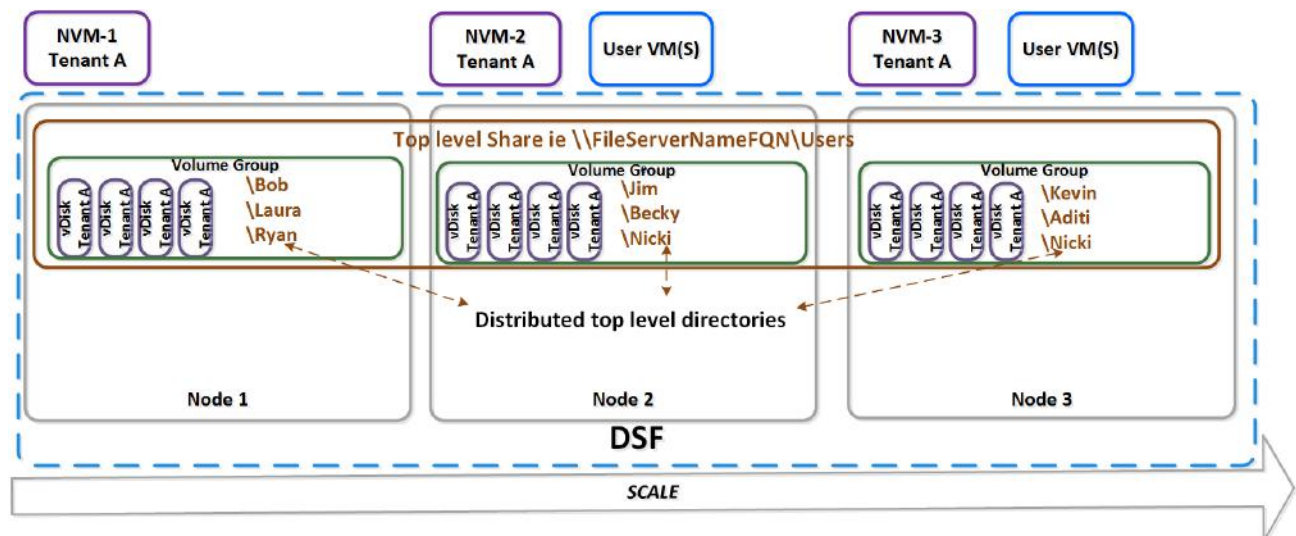


Figure 10: Distributed Directory Shares

Distributed share directories work well for home shares and exports because Nutanix Files automatically spreads the workload over multiple FSVMs per user (see the previous figure). If a user creates a share called “\\FileServer1\Users” that contains the top-level directories \Bob, \Becky, and \Kevin, \Bob may be on FSVM-1, \Becky on FSVM-2, \Kevin on FSVM-3, and so on. The FSVMs use a string hashing algorithm based on the directory names to distribute the top-level directories.

This distribution can accommodate a large number of users or directories in a single share or export. The scaling limits of more traditional designs can force administrators to create multiple shares or exports in which, for example, one set of users whose last names begin with A through M run off one controller and users whose names begin with N through Z run off another controller. This design limitation leads to management overhead headaches and unnecessary Active Directory complexity. For these reasons, Nutanix Files expects to have one SMB distributed share for the entire cluster. If you need to have more than one home directory share, you can create additional shares as needed.

The top-level directories act as a reparse point—essentially a shortcut. Consequently, administrators must create directories at the root of the share for optimal load balancing. We recommend that you set permissions at the share or export root before you deploy user folders. This step allows newly created top-level directories to inherit permissions, so you don't have to adjust them after the fact using the Nutanix Files Microsoft Management Console (MMC) plug-in.

Standard shares and exports don't distribute top-level directories. A single file server always owns the files and subfolders for standard shares and exports. The following diagram illustrates two standard shares (for example, accounting and IT) on the same file server.

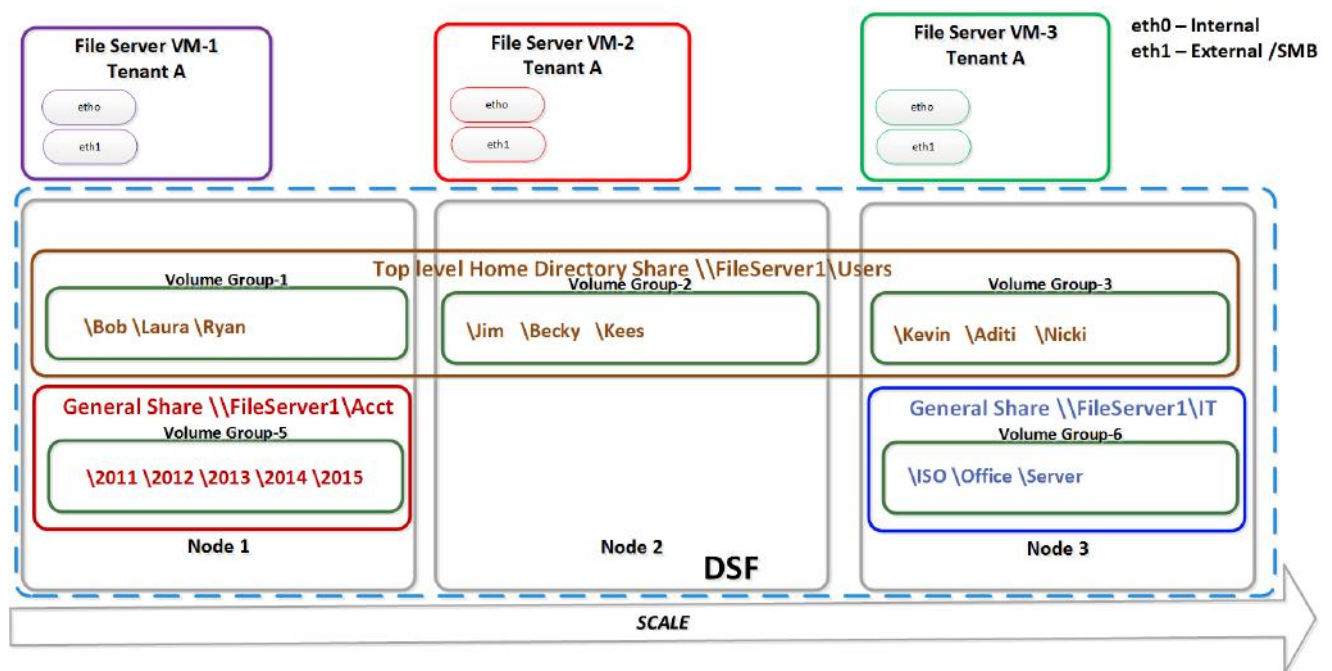


Figure 11: Two Standard Shares on the Same File Server

Standard shares and exports can store files in the root of the directory.

Managing NFS Distributed Share Directories

Distributed share directories with NFSv4 introduce some unique behaviors. To balance performance across FSVMs, each top-level directory you create becomes an automatically generated export. Nutanix Files mounts the export on demand when someone accesses the directory. Because these are exports rather than standard directories, it takes a few steps to remove a top-level directory. In this section we demonstrate how distributed top-level directories behave and then walk you through the process of deleting a top-level directory.

If you create a distributed share and mount it as `/projects`, you can see that mount on Linux using the `df` command:

```
# df /projects
Filesystem 1K-blocks Used Available Use% Mounted on
1.1.1.10:/projects 1073741824 839455744 234286080 79% /projects
```

Next, create some project directories and access them:

```
mkdir /projects/project1
mkdir /projects/project2
ls /projects/project1
ls /projects/project2
```

Accessing the directory, using ls in this case, mounts the automatically created top-level directory export. If you run df again you see two additional mount points:

```
# df | grep project
1.1.1.10:/projects 1073741824 839455744 234286080 79% /projects
1.1.1.11:/projects/project1 1073741824 839455744 234286080 79% /projects/project1
1.1.1.12:/projects/project2 1073741824 839455744 234286080 79% /projects/project2
```

These additional mount points allow a different FSVM to serve each export.

This behavior introduces additional steps when you delete a top-level directory. You can delete the project2 directory from any NFS client with the export mounted if you sign on as a user with the appropriate permissions. There are three steps to deleting a distributed share export:

- Delete the contents of the share:

```
rm -rf /projects/project2/*
```

- Unmount the project2 share:

```
umount /projects/project2
```

- Delete the project2 directory:

```
rmdir /projects/project2
```

At this point you've deleted the top-level directory and it becomes inaccessible to other clients mounting the export. Processes that access the export after you delete it receive a Stale File Handle error.

Distributed shares with NFSv3 clients behave differently than those with NFSv4 clients. NFSv3 connections don't automatically mount the top-level directories of distributed shares. Because the top-level directories aren't submounted, you can rename and delete them without the additional steps required for NFSv4.

Nested Shares

Nutanix Files 3.2 introduced support for nested shares. A nested share allows you to create a folder within an existing standard or distributed share and turn that folder into a directly accessible share. Both SMB and NFS protocols support nested shares. Nested shares inherit some attributes from the parent share, but you can modify other attributes. Nested share inherited attributes:

- Protocol type
- Maximum size
- Self-service restore
- Quota policy

Nested share modifiable attributes:

- SMB access-based enumeration (ABE)
- NFS authentication and access
- NFS advanced settings

Create a share/export ? X

Basics Settings Summary

Name
NestedShare

Description (Optional)

File Server
ADPAFS

Share Path (Optional) ?
/home1/tld1/newnestedshare

Cancel Next

Figure 12: Create Share Menu with Nested Share Path

Connected Shares

The connected shares function allows you to mount shares and exports as subfolders within other shares and exports. You can create a folder at any level of the folder hierarchy and mount either a standard or distributed share into that folder. Starting with Files 3.8, the share used as the parent hosting the folder can be a standard, distributed, or nested share.

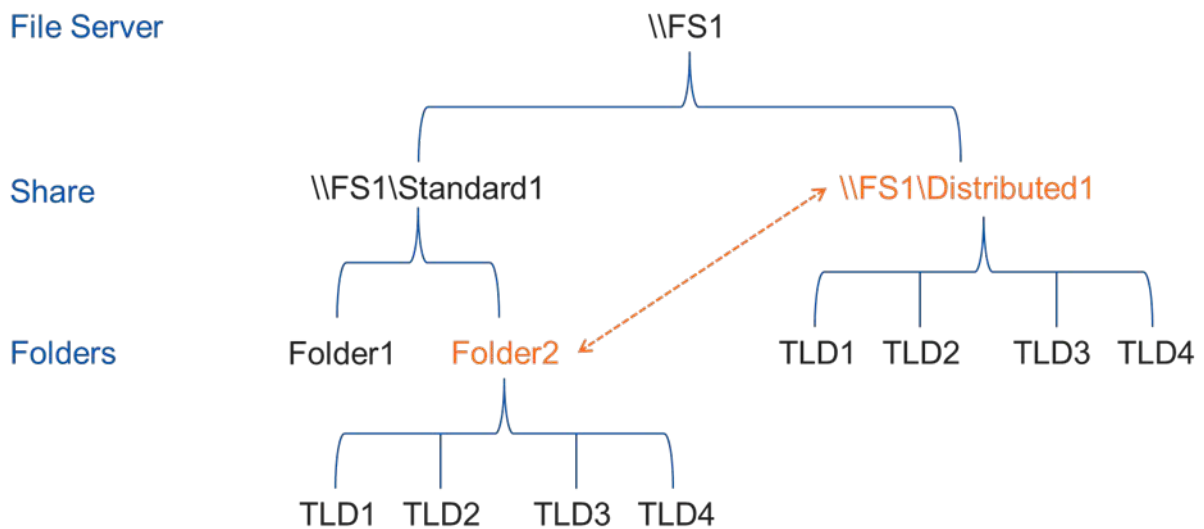


Figure 13: Submounting a Distributed Share into a Standard Share

There are several benefits of submounting shares and exports:

- You can use a distributed share at any level of a directory structure. Some applications have the folder paths where they store data or create directories hardcoded. If you submount, you can establish a distributed share at the level that matches these application requirements.
- A folder adopts the size limit of the submounted shares, which enables folder-level quotas.
- A folder adopts the self-service restore setting for a submounted share, which enables folder-level snapshots.

vDisk and Volume Group Allocation

Distributed shares and exports begin with 5 volume groups for each FSVM (for example, 15 for a three-node cluster). Nutanix Files distributes the volume groups to different FSVMs in the file server cluster. When you have a large number of users, multiple volume groups improve load balancing across the FSVMs.

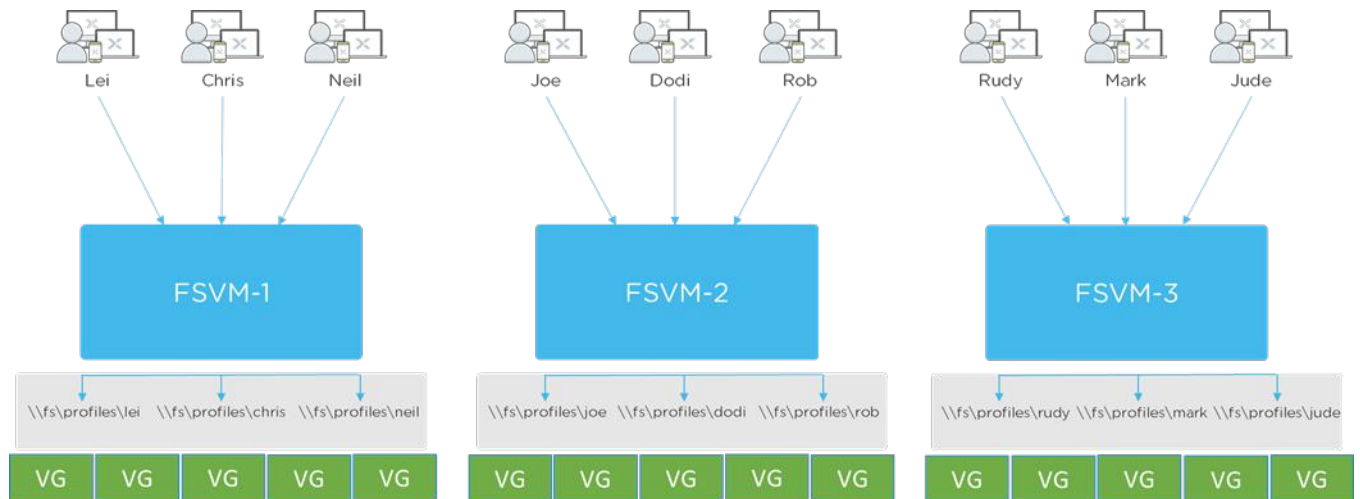


Figure 14: Distributed Share and Top-Level Directories

A single FSVM and volume group serve each standard share or export.

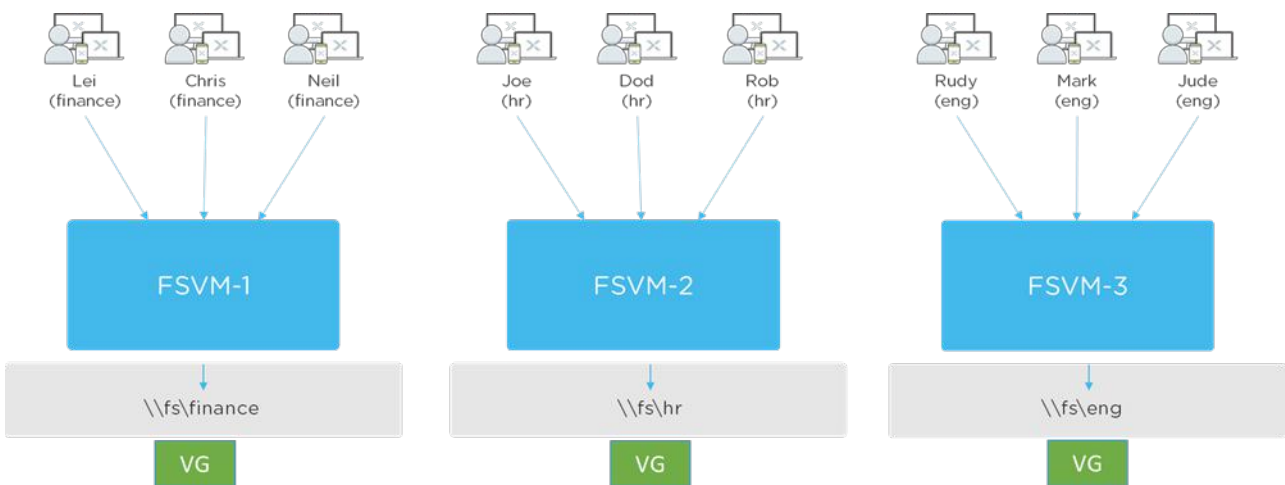


Figure 15: Three Standard Shares

Once you reach the limit of 10 volume groups per FSVM, new shares and exports use existing volume groups of the same type. For example, if you deploy a distributed share (15 volume groups) and 15 standard shares (creating 15 additional volume groups) on a three-node physical cluster, each FSVM hosts 10 volume groups: 5 for the distributed share and 5 for the standard shares. In this situation, because each FSVM is serving the maximum of 10 volume groups, the next share created uses an existing volume group.

Every file server maps one-to-one to a container. This mapping allows you to manage compression and erasure coding individually for each file server deployed. Inline compression and erasure coding are turned on by default to save capacity. We don't recommend turning on deduplication.

Tip: Use distributed shares for most of your use cases to ensure that user connections and data ownership are distributed across the FSVMs in the cluster.

From Files 3.6 onward, you can control compression on a share-by-share basis. When you create a share, compression is enabled by default.

Directory Layout

Nutanix Files can store millions of files in a single share and billions of files across a multinode cluster with multiple shares. To achieve good response times for environments with high file and directory counts, give some thought to directory design. Placing millions of files or directories in a single directory slows file enumeration, which may impact some applications.

For performance-sensitive applications, Nutanix recommends that you limit directory width—the number of files or directories in the root of a folder—to 100,000 objects. Increasing FSVM memory to cache metadata can help improve performance for environments with high object counts where directory enumeration is common.

Nutanix also recommends that you limit the number of top-level directories in distributed shares. The recommended number of top-level directories depends on the memory assigned to the FSVMs in the cluster. Your directory count shouldn't exceed 3,000 times the FSVM memory of one node. For example, for a three-node cluster with 12 GB of memory assigned to each FSVM, the maximum top-level directory count is $3,000 \times 12 = 36,000$.

Load Balancing and Scaling

Initial load balancing for a Files cluster is based on the number of standard shares or the number of top-level directories you have in a distributed share. As explained in the Exports and Shares section, a standard share resides on a single volume group owned by one FSVM at a time. A distributed share is a collection of volume groups containing top-level directories. If some standard

shares or top-level directories encounter a large workload and the volume groups supporting those workloads share an FSVM, a performance bottleneck can occur. Nutanix Files helps load balance workloads automatically when it detects potential performance problems by redistributing volume groups to different FSVMs for better load balancing across nodes.

Load balancing may occur in the following situations:

1. When an administrator removes an FSVM from the cluster.
2. When the distribution shares or top-level directories become poorly balanced during normal operation because of changing client usage patterns or suboptimal initial placement.
3. When increased user demand necessitates adding a new FSVM and its volume groups are initially empty.

Nutanix Files addresses the second and third situations by maintaining usage statistics and patterns to detect per-FSVM load (in terms of CPU and memory utilization) and per-volume group load (in terms of user connections and latency of operations). Nutanix Files uses these statistics to make a load balancing recommendation, but the administrator must accept the recommendation before Files carries out the action. Nutanix calls this feature one-click performance optimization.

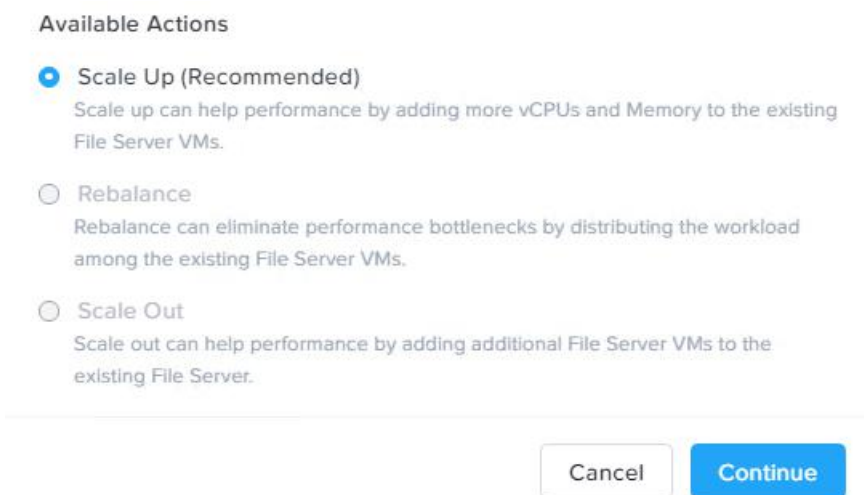


Figure 16: Performance Optimization Recommendation

Load balancing through volume group redistribution may not always improve performance. For example, if clients target a low-level share directory that can't be further distributed among FSVMs, performance doesn't improve. In such cases, Nutanix Files supports scaling up the FSVMs by adding vCPUs and memory. Scaling up is seamless to end users.

Note: We recommend a scale-up operation for performance optimization if SMB connection limits reach 95 percent utilization over a two-hour time window.

There is a brief outage during volume group migration and FSVM scale-out if you aren't using shares with continuous availability enabled. The file share or export requires a client reconnect after migration and scaling out. Today most clients try to reconnect for 50–60 seconds, which limits the overall impact.

Load balancing also occurs for each vDisk in a volume group. Nutanix Volumes requires the administrator to configure an iSCSI data service IP address. The data service IP address is a highly available virtual IP address used as an iSCSI discovery portal. Each vDisk in a volume group represents its own iSCSI target that any CVM in the cluster can own. Nutanix Volumes uses iSCSI redirection to place and automatically load balance these sessions as needed. The load balancing feature for Nutanix Volumes is called the Acropolis Dynamic Scheduler (ADS).

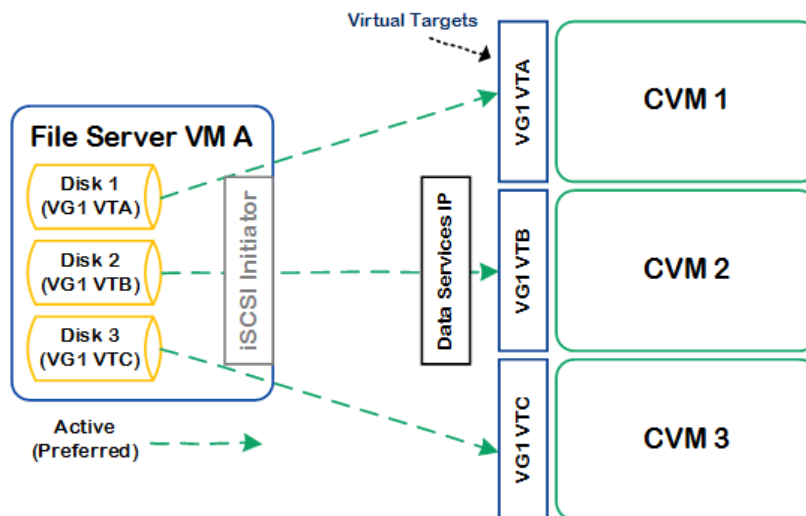


Figure 17: Load Balancing Volume Groups with Nutanix Volumes

For more detailed information, refer to the [Nutanix Volumes best practices guide](#).

High Availability

Nutanix designed Files to recover from a range of service disruptions, including when a local CVM or FSVM restarts or fails.

CVM Failure or Upgrade

If a CVM goes offline because of failure or planned maintenance, any active sessions against that CVM are disconnected, triggering the iSCSI client to sign on again. The new authentication occurs through the external data services IP, which redirects the session to a healthy CVM. The following figure shows this general process.

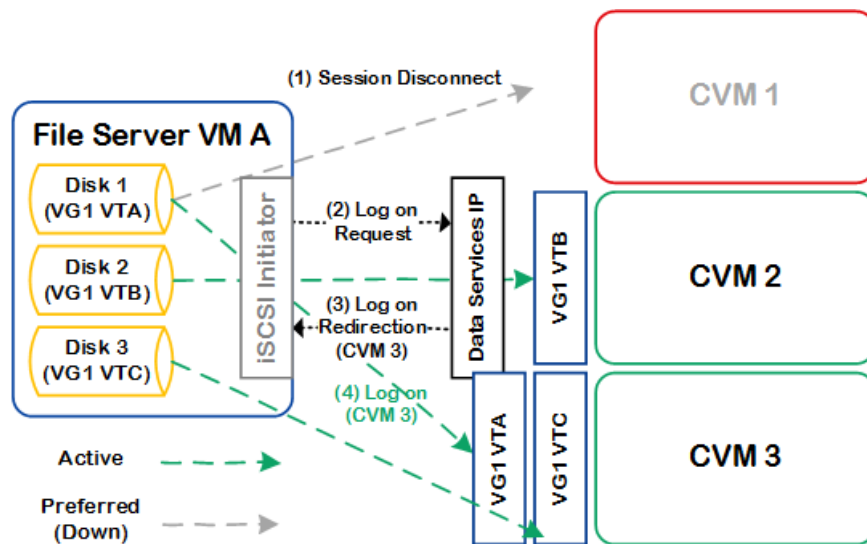


Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups

When the failed CVM returns to operation, the iSCSI session fails back. In the case of a failback, the system signs the FSVM off and redirects it to the appropriate CVM.

Node Failure

When a physical node fails completely, Nutanix Files uses leadership elections and the local Minerva CVM service to recover. The FSVM sends heartbeats to its local Minerva CVM service once per second, indicating its state and that it's alive. The Minerva CVM service keeps track of this information and can take action during a failover.

When an FSVM goes down, the Minerva CVM service unlocks the files from the downed FSVM and releases the external address from eth1. The downed FSVM's resources then appear on a running FSVM. The internal Zookeeper instances store this information so that they can send it to other FSVMs if necessary.

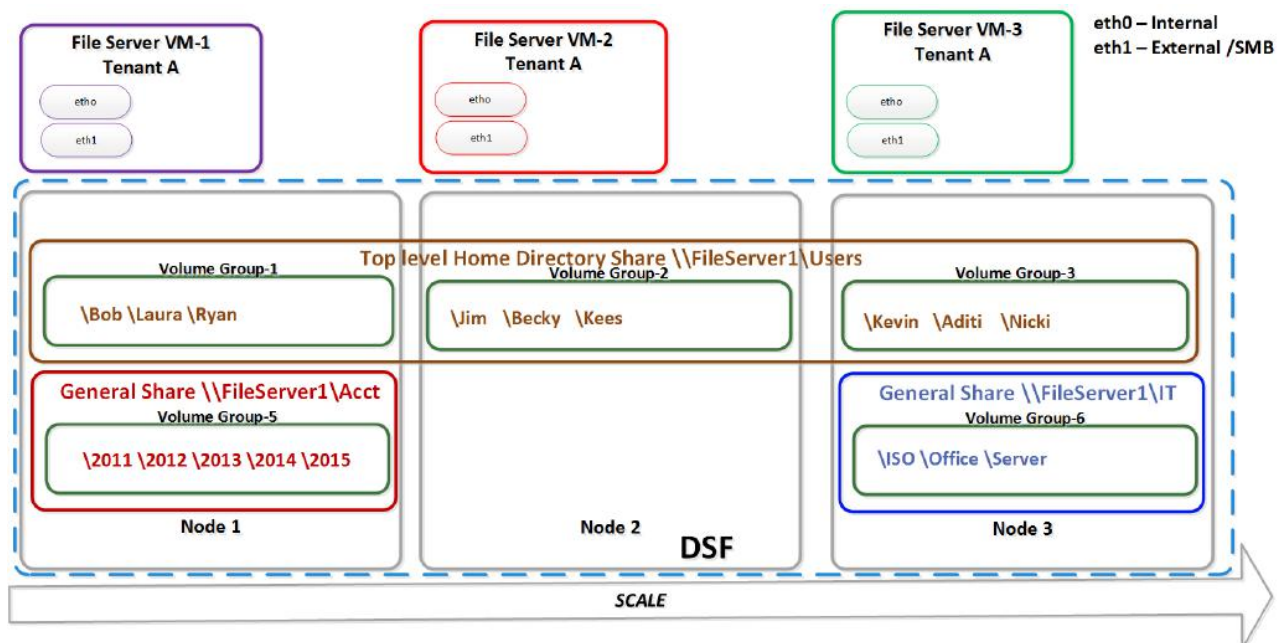


Figure 19: Each FSVM Controls Its Own Volume Groups

When an FSVM is unavailable, the remaining FSVMs volunteer for ownership of the shares and exports associated with the failed FSVM. The FSVM that takes ownership of the volume group informs the CVM that the volume group reservation has changed. If the FSVM that attempts to take control of the volume group is already the leader for a different volume group it has volunteered for, it relinquishes leadership for the new volume group

immediately. This arrangement ensures distribution of volume groups even if multiple FSVMs fail.

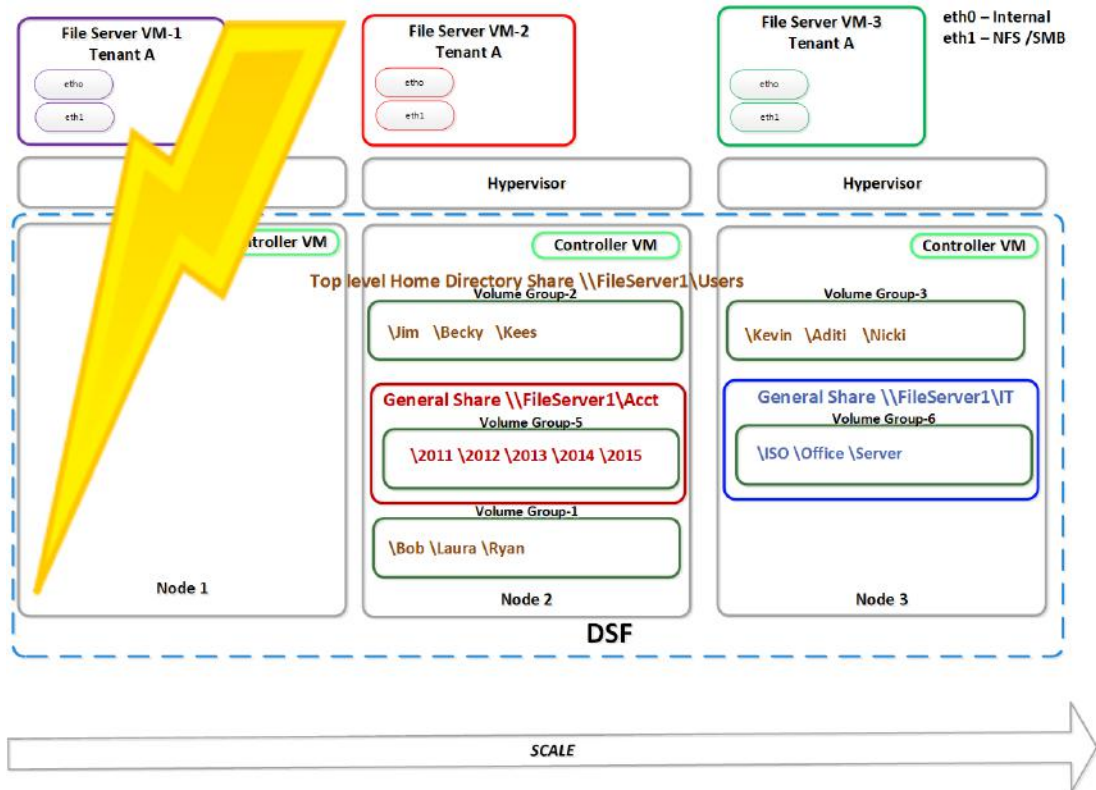


Figure 20: FSVM-1 Failure

The Nutanix Files Zookeeper instance tracks the original FSVM’s ownership using the storage IP address (eth0), which doesn’t float from node to node. Because FSVM-1’s client IP address from eth1 is now on FSVM-2, client connections persist. The volume group and its shares and exports are reregistered and locked to FSVM-2 until FSVM-1 can recover and a grace period has elapsed.

The failover process typically takes between 30 and 180 seconds, and client operations can be impacted during this time. SMB shares enabled for continuous availability and hard-mounted NFS shares experience I/O delays but maintain connections during high-availability events.

If an FSVM is offline for 30 minutes or more, the Nutanix Files cluster redistributes core cluster services (such as Zookeeper) owned by the offline node. The cluster needs at least three online FSVMs in order to move services. Moving cluster services allows the Files cluster to rebuild itself smaller and survive subsequent node failures.

When FSVM-1 comes back online and finds its shares and exports locked, it assumes that a high-availability event has occurred. After the grace period expires, FSVM-1 regains control of the volume group through the Minerva CVM service. Once the failed FSVM comes back online the cluster automatically switches back to the larger size.

To summarize, the process Nutanix Files goes through to reinstate control is as follows:

1. Stop SMB and NFS services.
2. Disconnect the volume group.
3. Release the IP address and share and export locks.
4. Register the volume group with FSVM-1.
5. Present new shares and exports to FSVM-1 with eth1.

Single-Node Environments

Nutanix Files clusters that only have one FSVM rely on hypervisor-based high availability to maintain services. If the physical node that owns the FSVM fails, the VM restarts on another physical node. Client connections are disrupted until the FSVM restarts.

Active Directory and SMB Operations

Nutanix Files SMB works with Active Directory. To deploy a cluster, you need domain privileges to create the machine account and the DNS entries for DFS referrals. The file server doesn't store these credentials.

Files 3.6 and later versions simplified the permissions required to add domains when you create a file server. The delegated domain user permissions required are as follows:

- Create computer objects.

- Read servicePrincipalName.
- Write servicePrincipalName.

Nutanix Files can also create the required DNS entries automatically during file server creation. Automated DNS entry creation requires Microsoft Windows DNS and a user account with DNS admin permissions.

The maximum number of client connections depends primarily on the amount of RAM in the FSVM. See the following table for Nutanix configuration recommendations as of the Files 3.7 release.

Table 2: Supported Active Client Connections

RAM per FSVM	Supported Client Connections per FSVM	Supported Client Connections for Four-FSVM File Server Cluster
12 GB	500	2,000
16 GB	1,000	4,000
24 GB	1,500	6,000
32 GB	2,000	8,000
40 GB	2,750	11,000
60 GB	3,250	13,000
96 GB	4,000	16,000

You can continue deploying additional FSVMs if you have free nodes; you can also deploy multiple file servers.

Nutanix Files uses DFS referrals to direct clients to the FSVM that owns the targeted share or top-level directory. The following diagram shows what happens behind the scenes when a client sends a file access request.

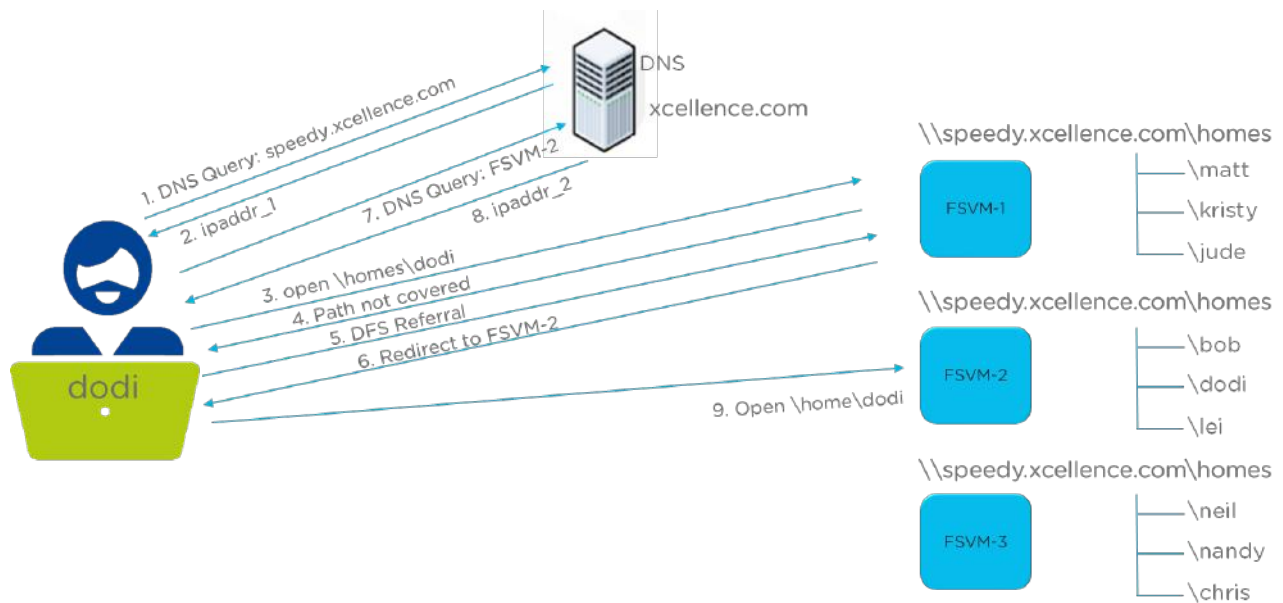


Figure 21: DNS Request for SMB

1. When the user Dodi accesses their files, they click on a shortcut that triggers a DNS request. The DNS request is first sent for the file server name.
2. Using DNS round robin, DNS replies with an FSVM IP address. In this example, the IP address for FSVM-1 returned first.
3. The client sends a create or open request to FSVM-1.
4. The `\homes\dodi` folder doesn't exist on this file server, so a `STATUS_PATH_NOT_COVERED` is returned.
5. The client then requests a DFS referral for the folder.
6. FSVM-1 looks up the correct mapping in the file server's Zookeeper and refers the client to FSVM-2.
7. A DNS request goes out to resolve FSVM-2.
8. The DNS request returns the IP address of FSVM-2.
9. The client gets access to the correct folder.

Managing Shares

You can manage Nutanix Files shares the same way you manage traditional file servers. For standard shares and top-level directories in distributed shares, you can assign permissions with native tools such as Windows Explorer. NTFS-

level permissions, also called Windows access control lists (ACLs) or NTACLs, manage all file and folder access.

The distributed share has some special requirements because of how we use DFS referrals. DFS referrals have a single namespace even though the data contained in the share may be spread out over many FSVMs. Typically, to delete a user's home directory, you select the top-level directory and delete the entire directory subtree. However, when using a distributed export or share you can't just delete the top-level directory because it's a separate share created internally by Nutanix Files. Because you can't directly delete a top-level directory, removing it requires additional steps.

There are two options for renaming or deleting distributed share folders:

1. Find out which FSVM is hosting the folder and rename or delete the folder directly using the FSVM as the UNC path.
2. Use the Nutanix Files MMC snap-in to manage top-level directories. Any file server administrator can perform the MMC operations; you don't need to assign privileges manually. Establish a connection to the Files namespace and perform the following SMB share management tasks:
 - a. Create, delete, rename, and change permissions for top-level directories.
 - b. Change NTFS permissions for shares.

Tip: Use the Nutanix Files MMC when you modify NTFS permissions at the root of a distributed share. The Files MMC ensures that permissions are propagated to the top-level directories based on inheritance.

You can modify share-level permissions with the Windows-native Shared Folders MMC. Simply launch the Shared Folders MMC and point it to the name of the Nutanix Files instance under the Another computer option. By default, all SMB shares have share-level ACLs set to Everyone with Full Control.

Note: We recommend leaving share permissions at Full Control and managing access with NTFS permissions.

You can also manage open file locks and user sessions with the Windows Shared Folders MMC.

Console1 - [Console Root\Shared Folders (\\FILESVR1)\Sessions]

File Action View Favorites Window Help

User	Comp...	Type	# Open Files	Connected Time	Idle Time	Guest
TME\fsctclient1u33	fsctclie...	Windows	0	00:00:10	00:00:00	No
TME\fsctclient1u77	fsctclie...	Windows	0	00:00:10	00:00:00	No
TME\fsctclient1u27	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u47	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u87	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u89	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u91	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u91	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u24	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u46	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u68	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u88	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u81	fsctclie...	Windows	0	00:00:13	00:00:00	No
TME\fsctclient1u17	fsctclie...	Windows	0	00:00:14	00:00:00	No
TME\fsctclient1u78	fsctclie...	Windows	0	00:00:14	00:00:00	No
TME\fsctclient1u85	fsctclie...	Windows	0	00:00:14	00:00:00	No
TME\fsctclient1u21	fsctclie...	Windows	0	00:00:17	00:00:00	No
TME\fsctclient1u5	fsctclie...	Windows	0	00:00:18	00:00:00	No
TME\fsctclient1u92	fsctclie...	Windows	0	00:00:18	00:00:00	No
TME\fsctclient1u26	fsctclie...	Windows	0	00:00:19	00:00:00	No

Figure 22: Shared Folders MMC

Access-Based Enumeration

Access-based enumeration (ABE) is a Windows feature (SMB protocol) that filters the list of available files and folders on the file server to include only those the requesting user can access. This filter helps both users and administrators, saving time for the user and worry for the administrator concerned about users accessing files not meant for them.

ABE doesn't control security permissions, and running ABE has associated overhead. Every time a user requests a browse operation, ABE must filter out objects the user doesn't have permission for. Even if the user has permission to access all contents of the share, ABE still runs, causing additional CPU cycles and increased latency.

Tip: Home directories (using a distributed share) are a great example of where you shouldn't enable ABE. As most users get their home mapping when they sign on and always have access to their own content, we don't recommend enabling ABE for home directory use cases.

SMB Signing

SMB signing is a feature of the SMB protocol that enables a digital signature against each network packet. Digital signing helps ensure the packet's origin and authenticity and prevent tampering, such as eavesdropping attacks against data in-flight.

Nutanix Files supports SMB signing and honors the signing request as configured from the SMB client. Nutanix Files doesn't require a configuration setting; just enable SMB signing from the client and Nutanix Files negotiates the appropriate setting.

The SMB protocol uses AES-CMAC to compute signatures for SMB signing. This computation process can impact performance against the SMB share where you enabled signing. With Nutanix Files 3.2 and later versions, the signature computation uses the Intel processor AES-NI instruction set. The Intel processor hardware acceleration helps reduce the overhead associated with SMB signing.

Note: SMB signing lowers the maximum performance possible for SMB shares.

SMB Encryption

Nutanix Files 3.6 introduced support for in-flight encryption of SMB data between clients and shares. SMB encryption is disabled by default with Nutanix Files. You can enable encryption as needed on a share-by-share basis.

Create a share/export ? x

Basics **Settings** Summary

Use "Distributed" share/export type instead of "Standard" ?

Enable Self Service Restore ?

Enable Access Based Enumeration (ABE) ?

Blocked File Types ?
You can also block file types on the file server(all shares) from file server update

Encrypt SMB3 Messages ?

Figure 23: Enable SMB Encryption

Clients must support encryption in order to access shares with encryption enabled. Nutanix Files offloads encryption processing using the Intel processor AES-NI instruction set to limit performance impact.

Note: SMB encryption lowers the maximum performance possible for SMB shares.

Durable SMB File Handles

Durable handles let SMB clients survive a temporary client-side connection loss after opening a file, allowing transparent client reconnection within a timeout. Starting with the 3.6 release, Nutanix Files supports durable handles for SMB clients with no configuration required. SMB 2.x and SMB 3.0 clients can use durable handles to reconnect transparently if there's a client-side network interruption. The system doesn't maintain durable handles through file server events such as FSVM failures or upgrades.

SMB 3.0 Transparent Failover

SMB 3.0 Transparent Failover is a feature of the SMB 3.0 protocol that enables fully nondisruptive operations against SMB shares. Transparent failover is often called continuously available file shares. File server-side failures or upgrade events persist file handles so that clients transparently reconnect to another

FSVM without impacting applications. Nutanix Files versions 3.7.1 and later support SMB 3.0 Transparent Failover. Continuously available shares are intended for specific use cases, like Citrix App Layering and FSLogix, that require nondisruptive operations.

Note: SMB Transparent Failover enforces synchronous write operations, which may decrease the maximum possible performance for SMB shares.

DFS Namespace

DFS Namespaces (DFS-Ns) are commonly used to logically organize shared folders on different servers or in different geographic locations into a single namespace. These shared folders appear to the user as a unified hierarchical directory they can navigate using any Windows client. The server names and their locations are completely hidden from the user, enabling large scale-out architectures.

The most common architecture is an Active Directory-integrated DFS-N, where the namespace is hosted on two or more domain controllers (namespace servers) and the file data is stored on member servers. Nutanix Files supports the use of DFS-N when you use Nutanix Files as a member server and Nutanix Files shares as folder targets in the namespace. When you use Files 3.5.1 or later versions, you can use either distributed shares or standard shares.

DFS-N is also commonly used in active-active replication scenarios. DFS-N allows multiple file servers hosting the same data to support a common folder. DFS-N helps provide site affinity, based on Active Directory, for users to connect to the file server acting as a folder target. You need a replication engine to maintain the data when active-active scenarios are present. Nutanix supports the use of Peer Software as the replication engine for active-active scenarios. To learn more about the Peer Software and Nutanix integration, see [the PeerGFS and Nutanix Files datasheet](#).

Active Directory, LDAP, and NFS Operations

NFSv4

Nutanix Files supports NFSv4, which is more stateful than NFSv3 and includes advanced features such as strongly mandated security and DFS-like referrals.

Moreover, most recent distributions of Linux, Solaris, and AIX use NFSv4 as the default client protocol.

Nutanix Files supports Active Directory, LDAP, and unmanaged access to NFSv4 exports. To make the transition from NFSv3 easier, Files doesn't require administrators to configure Active Directory or LDAP. You can use AUTH_SYS or AUTH_NONE authentication. AUTH_SYS authenticates at the client, just like NFSv3.

There are three different levels of Kerberos authentication when you enable Active Directory support. Each of the following options uses Kerberos version 5:

1. krb5, DES symmetric key encryption, and an MD5 one-way hash for Nutanix Files credentials.
2. krb5i, in addition to krb5, uses MD5-based MAC on every request and response.
3. krb5p, on top of krb5 and krb5i, makes the connection between client and server private by applying DES encryption.

Directory Services

Select the protocols you plan to use and configure directory services options.

Use SMB Protocol

Active Directory is necessary for using SMB protocol

Active Directory Realm Name [Leave Domain](#)

tme.local

Use NFS Protocol

You can use AD, LDAP or Unmanaged for NFS user management and authentication

User Management And Authentication ⓘ

Active Directory

Enable Identity Management for Unix (RFC 2307) ⓘ

Active Directory Realm Name: User account location to join

tme.local

Cancel Update

Figure 24: Nutanix Files User Management with NFS

When you deploy Nutanix Files for NFS, you can select Active Directory or LDAP or leave it unmanaged. The following diagram shows what happens behind the scenes when a client sends a file access request using NFS.

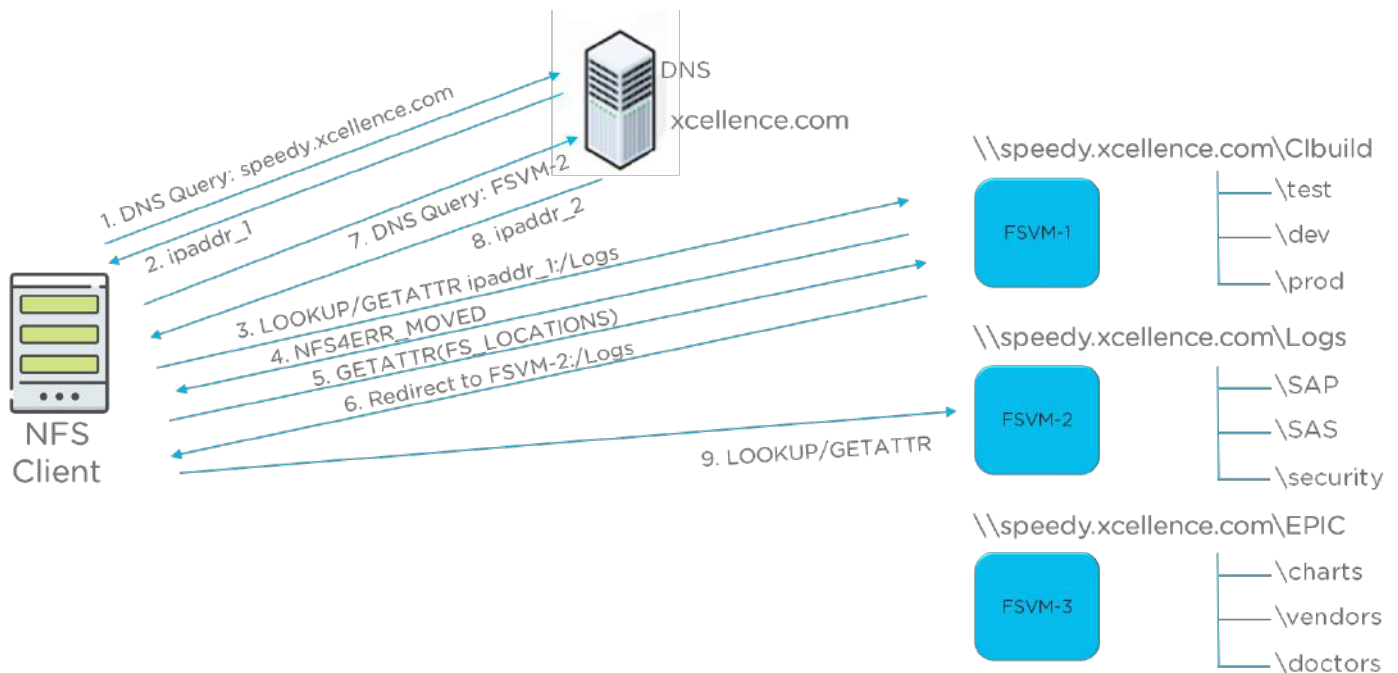


Figure 25: DNS Request for NFSv4

1. When the server wants to access files, the client first sends a DNS request for the file server name.
2. Using DNS round robin, a DNS reply returns with an FSVM address. In this example, the IP address for FSVM-1 returned first.
3. The client sends a create/open request to FSVM-1.
4. The \Logs mount doesn't exist on this file server, so it returns NFS4ERR_MOVED.
5. The client then requests a GETATTR(FS_LOCATIONS).
6. FSVM-1 looks up the correct mapping in the file server's Zookeeper and refers the client to FSVM-2.
7. A DNS request goes out to resolve FSVM-2.
8. The DNS request returns the IP address of FSVM-2.
9. The client gets access to the correct mount point.

NFSv3

Nutanix Files supports NFSv3. Each file server has NFSv3 enabled by default, but you can manually disable it. All NFS exports have NFSv3 enabled or disabled based on this file server setting. Files 3.5 supports LDAP and unmanaged exports with NFSv3 but doesn't support Active Directory and Kerberos for NFSv3.

Note: You must mount exports with the TCP protocol if clients use NFSv3.

Nutanix Files with NFSv3 includes support for both distributed and standard exports. Unlike NFSv4, NFSv3 doesn't support DFS-like referrals. Clients aren't redirected to the FSVMs that host a given distributed top-level directory or standard export; they connect to the first FSVM resolved using DNS. An internal remote procedure call (RPC) manages any file access requests to exports and top-level directories not hosted by the client-connected FSVM. This RPC sends or receives data between the client-connected FSVM and the FSVM that owns the required exports and volume groups.

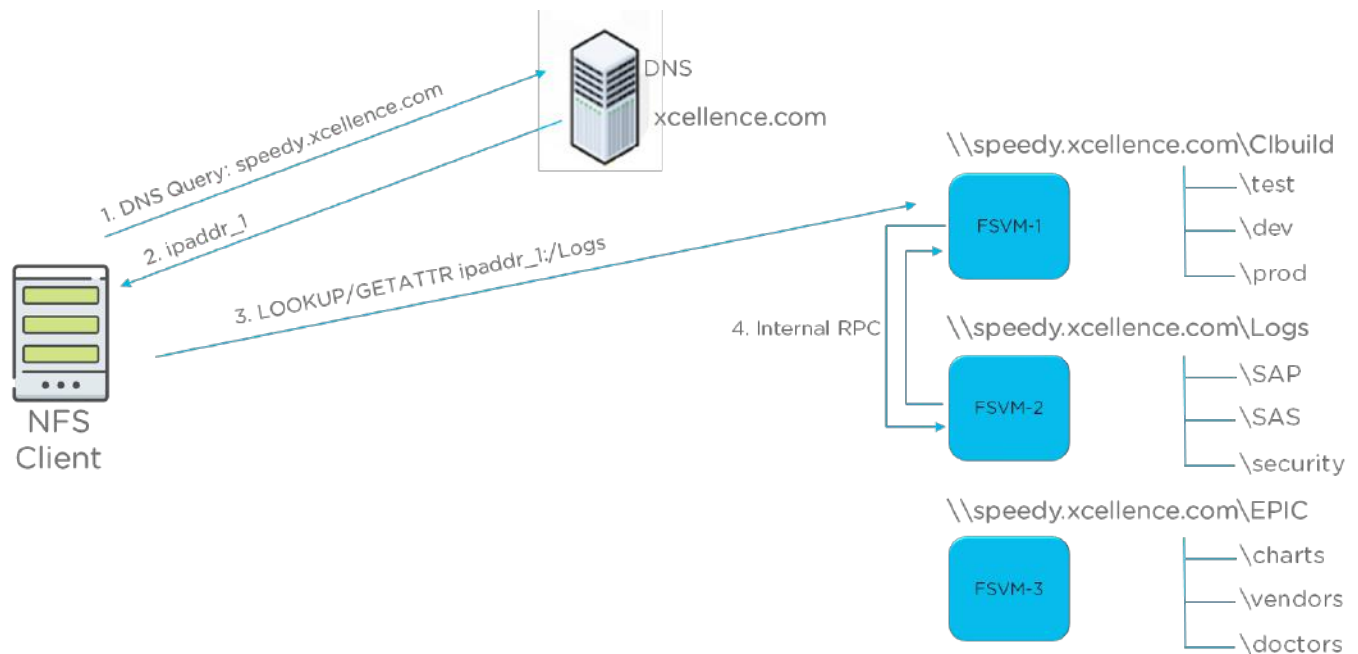


Figure 26: DNS Request for NFSv3

Multiprotocol

From the Nutanix Files 3.5.1 release onward, you can create file shares that are accessible from both SMB and NFS clients, referred to as multiprotocol shares. Multiprotocol shares allow simultaneous read access to the underlying file data from either protocol. Write access can also occur from either protocol, but not simultaneously to the same file. Authentication support for multiprotocol shares includes all protocols and authentication options available with Nutanix Files (such as Active Directory for SMB and LDAP, AUTH_SYS, AUTH_NONE, and Active Directory support for NFS).

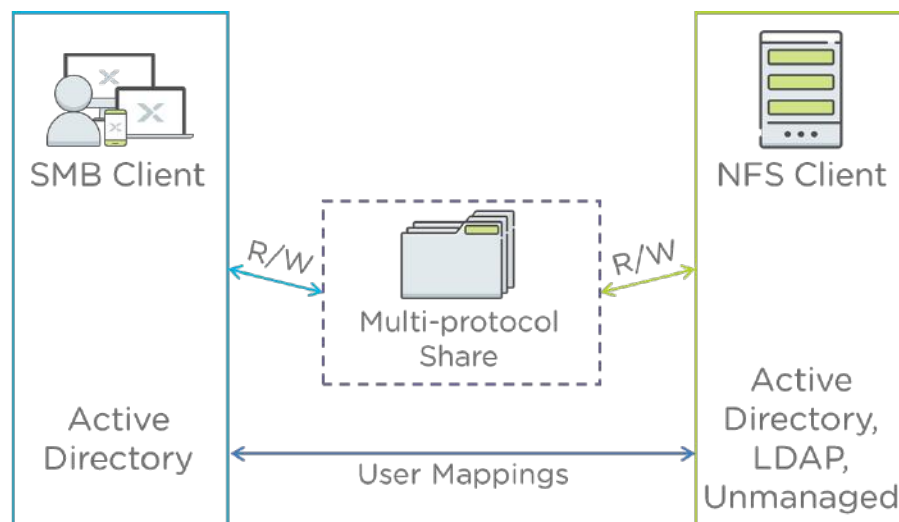


Figure 27: Multiprotocol Shares

Multiprotocol introduces the concept of a native protocol (either SMB or NFS) for a given file share. You specify the native protocol and non-native protocol when you create the share.

? X

Basics Settings Summary

Description (Optional)

Multi-protocol Native SMB

File Server

filesvr1

Share Path (Optional) ?

<sharename>/<dir>

Max Size (Optional)

GiB

Select Protocol

NFS

SMB

Also enable non-native NFS Access

Cancel Next

Figure 28: Enable Non-Native Protocol Access

You manage all access control for a multiprotocol share using the native protocol. For SMB, the native protocol is Windows ACLs, and for NFS, it's Unix mode bits. When non-native protocol access occurs, Nutanix Files maps user access to the permission applied with the native protocol. A mapping must exist between the user accounts using the non-native protocol and the user accounts with permission applied through the native protocol.

If you use Active Directory for both SMB and NFS with Kerberos, you don't need any explicit mappings because both have the same users and groups. You can use Identity Management for Unix (RFC 2307) to map NFS users to your Active Directory users with attributes. For LDAP or unmanaged accounts, or for Active Directory without Kerberos, you must create a mapping between the users and groups in Active Directory. You can use Prism or the nCLI to manage Nutanix Files user mappings.

Figure 29: Multiprotocol User Mapping

You can configure a default index to map all non-native users and groups to a specific native user or group. You can also use a rule-based mapping for Active Directory and LDAP users—specifically, a template where the SMB name matches the NFS username. Additionally, you can configure explicit mapping, which overrides rule-based mapping. Explicit mapping consists of two mapping subcategories: one-to-one mapping lists and wildcard mapping. You can use the one-to-one mapping list to manually enter or upload a CSV file that maps users or groups across protocols. Use wildcards for many-to-one mapping. You can also deny share access for a specific user or group. For more details on how to create user mappings, see the [user mapping section of the Nutanix Files guide](#). The following table shows the user mapping requirements depending on the chosen directory service and authentication type.

Table 3: User Mapping Requirements

SMB Directory Service	NFS Directory Service	Export Authentication Type	Supported	User Mapping Required
AD	AD	Kerberos5*	Yes	No
AD	AD	System	Yes	Yes (Name to ID mapping)

SMB Directory Service	NFS Directory Service	Export Authentication Type	Supported	User Mapping Required
AD	AD	None	Yes (Primary NFS only)	Yes (Name to ID mapping)
AD	AD + RFC2307	Kerberos5*	Yes	No
AD	AD + RFC2307	System	Yes	Yes (Name to ID mapping)
AD	AD + RFC2307	None	Yes (Primary NFS only)	Yes (Name to ID mapping)
AD	LDAP	System	Yes	Yes (Name to Name mapping)
AD	LDAP	None	Yes (Primary NFS only)	Yes (Name to Name mapping)
AD	Unmanaged	System	Yes	Yes (Name to ID mapping)
AD	Unmanaged	None	Yes (Primary NFS only)	Yes (Name to ID mapping)

* Kerberos 5, 5i, and 5p

Quotas

Administrators can configure the default, user, and group quota space for any share. The default level is the quota limit for every user unless the administrator specifies otherwise. A user-level quota policy sets a specific amount of storage for a single user. For example, if an administrator allocates 1 GB, the user can't take more than 1 GB. A group-level quota policy extends a user policy to include all users for an entire Active Directory group, where each user can use the assigned quota value. For example, if the administrator sets a group's quota to 10 GB, then each member of that group can use 10 GB. You can assign quotas for SMB, NFS, or multiprotocol-enabled shares. In the case of multiprotocol shares, administrators apply quotas using the native protocol. User or group policies are enforced for the non-native protocol based on the user mapping.

Quota Policy ? X

Add a quota policy
Add a quota policy for a specific user or each user in a group.

User

Group (Quota will be applied to each member of group individually)

USER OR GROUP

domain users

QUOTA

20 GiB

ENFORCEMENT TYPE

Hard Limit
User will be put in read only mode and will receive email notification daily until remediation.

Soft Limit
User will receive daily email notification but will not be put in read only mode.

Alert Emails
Quota alert emails are sent at 90% and 100% of quota limit.

Cancel Save

Figure 30: Setting a Share Quota

Notifications

Administrators can configure Prism to send email alerts to the user and to other recipients using the same engine that sends cluster alerts. Designated users receive email notifications when the quota is near maximum consumption—a warning email at 90 percent and an alert email at 100 percent. You can also add departmental share owners to the email notification list so they know they may need to take action. With the Files 4.0 release you can customize the email template used for notification.

The following table shows the order of precedence when dealing with quotas for a share.

Table 4: Order of Precedence for Quotas

Order of Precedence	Policy
1	User policy
2	Group policy (the group policy with the highest quota wins)
3	Default user policy

Enforcement

The administrator can also configure enforcement types for each quota-level category. Enforcement types determine if a user or group can continue to use the share after they've consumed their quota. A hard enforcement type prevents the user from writing on the share when they reach their quota limit. A soft enforcement type allows a user to write even if they exceed the quota limit. Under either enforcement type, users over their quota receive an email notification every 24 hours until they resolve the issue.

Selective File Blocking

You can define a list of file names and file extensions to block from storage on SMB, NFS, and multiprotocol-enabled shares. You can define a list of files to block at the server level and at the share level. These entries can include wildcards for both the file name and file extension. For example, you can block the file pattern `encrypt *.* xt`.

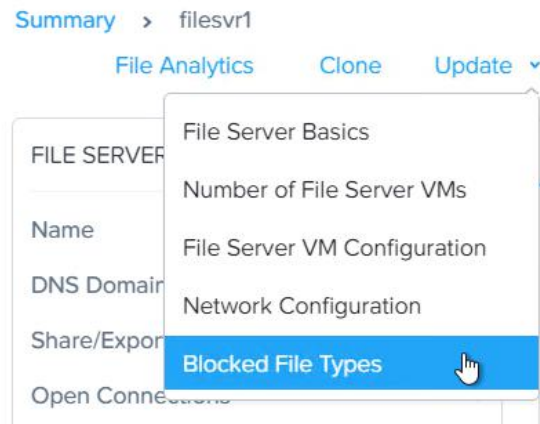


Figure 31: Blocked File Types for the File Server

A comma-separated list of file names and extensions defines the blocked file types. File types blocked at the server level apply to all shares. When you define a list of file types at the share level, the share-level setting overrides the server-level setting.

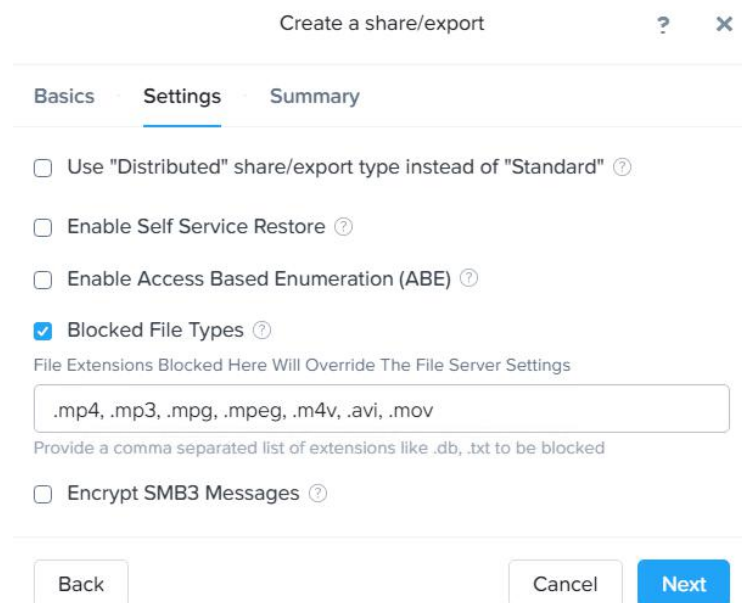


Figure 32: Blocked File Types for a Share

When you attempt to create a file with or rename a file to a blocked pattern, an access denied message appears. You can read, edit, or delete existing files

created prior to a blocked file type policy. You can also use File Analytics to discover if any unwanted file types exist on a file server before you apply the policy. With version 3.0 or later, File Analytics can also apply a list of file patterns matching known ransomware variants.

Note: When you use Data Lens to manage your ransomware file name patterns, you can't see the patterns from within the file server or share specific blocking lists.

File Analytics

Nutanix Files offers a rich auditing API to which applications can subscribe and receive real-time notifications of file-related events, including file creation, deletion, read, write, and permissions changes. A common use for these APIs is forwarding such events to a syslog server for retention and audit trails. While logging audit trails is useful, we needed to simplify insights into this data, so we developed Nutanix File Analytics to consume this native auditing API while delivering additional insights to the underlying data and user activity.

You can deploy File Analytics using Prism: simply download the Analytics bundle and perform a one-click deployment. Once deployed, launch the File Analytics page and enable analytics against the desired file server instances running on your cluster. Enabling analytics requires a user account with administrative permissions to do an initial scan of the file server.

Intelligent Insights

File Analytics scans and then continuously captures all file activity for registered file server instances. This logging helps form a repository of information so administrators can review what operations have occurred against specific data and by specific users. File Analytics analyzes logged events to provide an initial dashboard of information:

- Capacity trend, which shows what's being consumed and how it's changed over time.
- Data age, which is the calculation of the last time a file was modified and the percentage of data at varying age ranges.
- Anomaly alerts, which show all file operations that exceed a given anomaly threshold, like the deletion of many files.

- Permission denials, which are the number of permission denial events for specific users over the selected time range.
- File distribution by size, which shows the number of files in a given size range.
- File distribution by type, which details storage consumption by file type, such as log files, image files, and video files.
- The top five active users based on total operations over the selected time period.
- The top five accessed files based on total operations over the selected time period.
- File operations, which detail the most frequent types of operations, like file create, read, or write, over a selected time period, including the trend over that time.

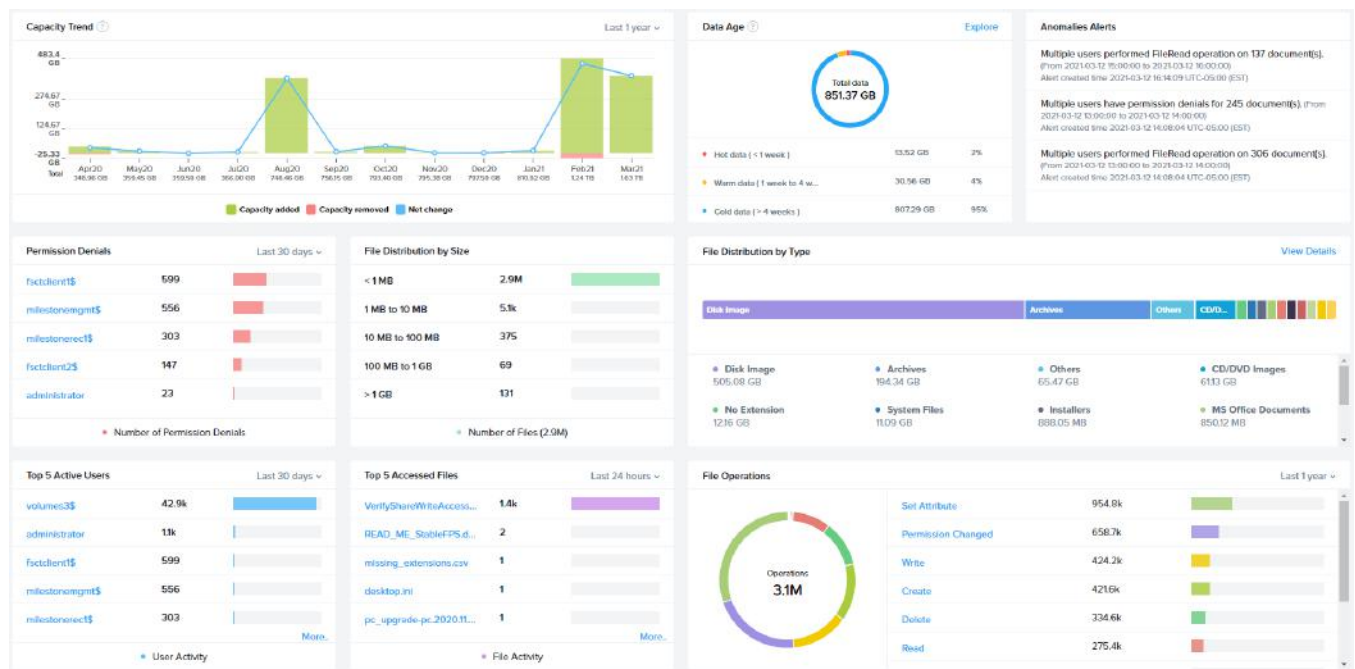


Figure 33: File Analytics Dashboard

The dashboard provides a quick and easy health check so you better understand capacity trends, data age, and file activity that could indicate malicious activity.

Data Age Analytics

Starting with version 3.0, File Analytics added an extension called data age analytics to the data age widget on the dashboard. Data age analytics shows you how frequently users access your data over time ranges you define. Once you customize the data age ranges to match your requirements, data age analytics shows you the hot, warm, or cold data trends over the specified timeframe, such as the last week, month, or longer interval. You can also see the growth percentage of a given category in your designated range.



Figure 34: Data Age Analytics

Audit Trails

Audit trails allow you to search for a specific file or a specific end user to find all file or user activity for a given timeframe. You can search for a given file or user based on wildcards. The frequency and types of operations, by users and against files, including the time they occurred, are displayed over your specified time period. You can further filter the audit trail based on operation type, such as open, read, write, delete, and other events.

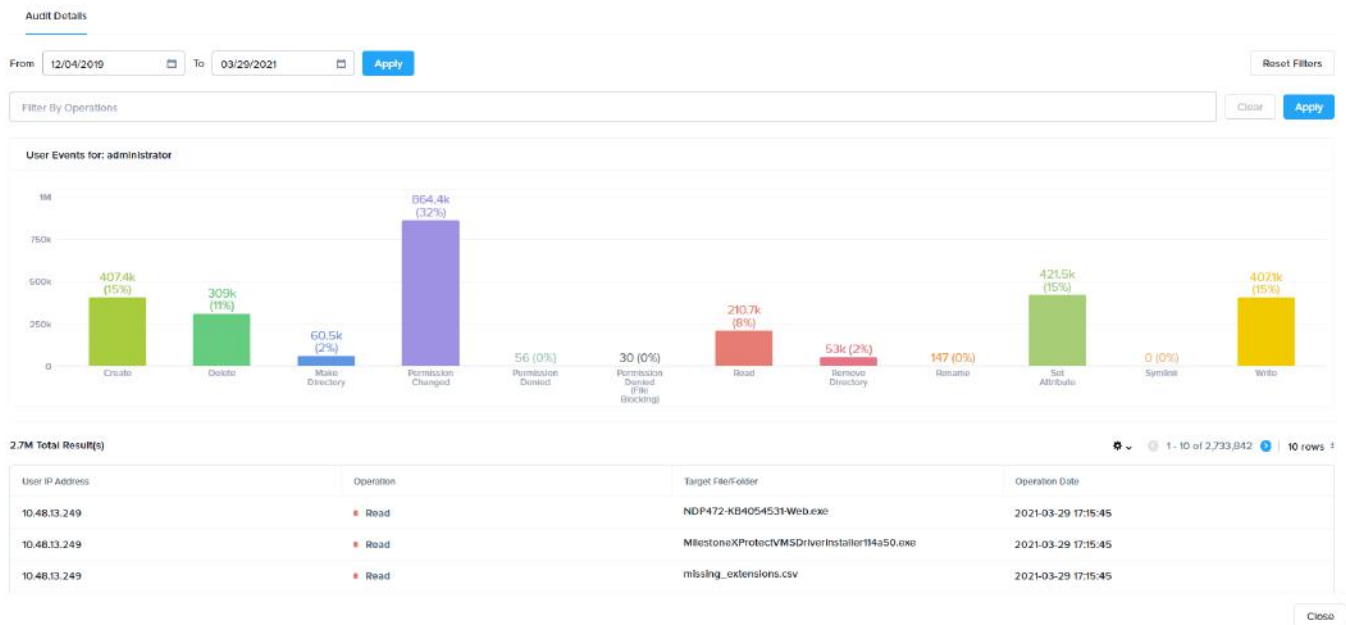


Figure 35: File Analytics Audit Trails

The ability to search user and file activity helps you monitor access activity against sensitive data. You can also download the queried activity to a JSON- or CSV-formatted file for further reporting activities.

Anomaly Detection

File Analytics allows you to define anomaly alerts that represent specific operations as run by an individual or against the file server as a whole.

Define Anomaly Rules ✕

Setup your anomaly policies and anomaly email recipients for fileserver "filesvr1" here.

Anomaly Email Recipients

Add one or more comma separated email addresses if you want to receive email alerts.

administrator@tme.local

[+ Configure new anomaly](#)

Events	Minimum Operati on % ?	Minimu m Operati on Count ?	User ?	Type ?	Interval ?	Actions
Permission Denied (File ...	3	500	All Users	Hourly	1	✎ ✖
Permission Changed	1	100	All Users	Hourly	1	✎ ✖
Permission Denied (Acc...	3	150	All Users	Hourly	1	✎ ✖
Read	1	100	All Users	Hourly	1	✎ ✖
Create	1	750	All Users	Hourly	1	✎ ✖

Cancel
Save

Figure 36: File Analytics Anomaly Detection

You can specify different events, such as permission changes, permission denials, file deletes, and file creates; define an operation percentage and operation count as a part of the alert; specify the interval for when these thresholds apply to the operations in question; and specify email recipients for the anomaly events.

Tip: Minimum operation percentage is based on the number of files. For example, if there are 1,000 files and the minimum operation percentage is 5, it means the defined event (such as create or delete) impacted 50 files within the specified interval.

As anomaly events occur, the email addresses defined in the anomaly rule receive detailed information. You can also monitor each anomaly alert and track it over time with the usage anomalies dashboard. The dashboard shows you the users who caused the anomalies and the folders the anomalies impacted and details the types of anomalies and trends.



Figure 37: File Analytics Usage Anomalies Dashboard

Strange access patterns that trigger anomalies can indicate malicious user, virus, or ransomware activity. Seeing this activity and being alerted in real time can help you stop these operations and better understand user access patterns across your organization.

Ransomware Intelligence

Ransomware is a persistent concern that requires multiple security controls and software layers to mitigate. Nutanix Files has long supported centralized antivirus scanning through ICAP (Internet Content Adaptation Protocol) with security vendors like Trend Micro, McAfee, BitDefender, and Symantec. Nutanix offers a comprehensive approach to ransomware across our portfolio of products that you can read about in more detail in our [Ransomware Threat solution brief](#). To seriously respond to ransomware, a solution should prevent the infiltration of ransomware and malware, detect any infection attempts, alert the organization and initiate defensive measures, and, if the worst happens, provide a comprehensive strategy for recovery.

Nutanix Files and File Analytics have many of the core features required to help detect, protect, analyze, and recover from ransomware. File Analytics 3.0 begins the journey to combine these technologies into a comprehensive interface, built so you can manage your ransomware strategy with Nutanix Files. It starts with a dedicated ransomware dashboard in File Analytics that summarizes any detected vulnerabilities, including the impacted shares and

clients that may be compromised. The dashboard also shows you whether your shares are protected with SSR snapshots. You can enable SSR against the unprotected shares from this interface.

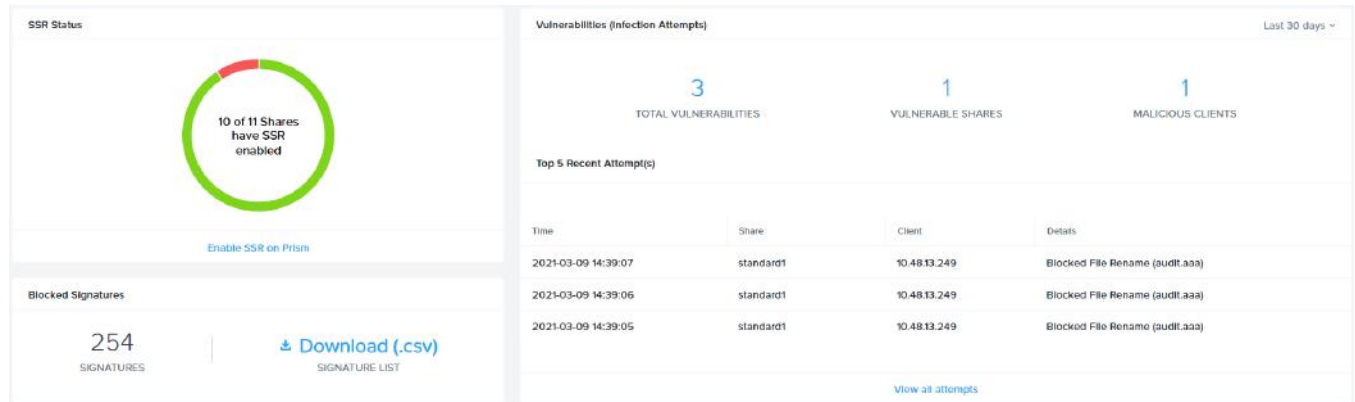


Figure 38: File Analytics Ransomware Page

Nutanix Files can block file creation or file rename operations for specific file extensions. Files 3.8 extended this feature to support wildcards with file names and file extensions. When you enable ransomware protection in File Analytics, Files automatically adds the names and extensions of known ransomware variants to the blocking list. If anyone attempts any file creation or file rename event involving these blocked file types, the event appears as a vulnerability that Files reports on the dashboard and emails to the designated users.

Custom Reporting

File Analytics captures real-time user audit data and file metadata for Nutanix Files environments. The 3.0 release of File Analytics enabled you to mine this data more effectively by introducing a custom reporting page, so you no longer needed to rely on the information available on the dashboard or through audit search.

When creating a new report, you first select an entity, such as files, folders, audit events, or users. You can then choose which attributes to filter with customizable values. For example, you can choose attributes like file size greater than or less than a given value, file age based on access time or creation date, or all audit events within a given time range. Finally, you can

choose which columns of data to include in the report, like user and file names, paths, clients, operation types, and other details associated with the entity.

The screenshot shows the 'Report Builder' interface with the following configuration:

- Step-1: Define Report Type:** Based on 'Files'.
- Step-2: Add/Remove column in this report:** Columns include 'file_path(File Path)', 'object_name(File Name)', and 'share_UUID(Share Name)'.
- Step-3: Define Filters:**
 - Filter 1: Attribute 'size', Operator 'greater_than', Value '500 MB'.
 - Filter 2: Attribute 'creation_date', Operator 'less_than', Value '09/01/2020'.
- Step-4: Define maximum number of rows in this report:** Count '100'.
- Sort by:** Attribute 'size'.

Buttons for 'Generate report' and 'Run Preview' are visible. Below the configuration is a 'Report Preview' table:

file_path(File Path)	object_name(File Name)	share_UUID(Share Name)
filesvr1/standard1/iso/GravityZoneEnterprise.raw/GravityZoneEnterprise...	GravityZoneEnterprise.raw	standard1
filesvr1/standard1/iso/Windows/en_windows_server_2016_updated_feb...	en_windows_server_2016_updated_feb_2018_x64_dvd_11636692.iso	standard1
filesvr1/standard1/iso/AOS/5.11.2/nutanix_installer_package-release-eup...	nutanix_installer_package-release-euphrates-5.11.2-stable-x86_64.tar.gz	standard1
filesvr1/standard1/iso/AOS/5.11.1/nutanix_installer_package-release-euph...	nutanix_installer_package-release-euphrates-5.11.1-stable-x86_64.tar.gz	standard1

Figure 39: Report Builder

Once you define the report, you can save it and run it again at any time.

The screenshot shows the 'Download Reports' section with a '+ Create a new report' button and a table of saved reports:

Name	Status	Last Run	Actions
Files >= 100GB	Success	2021-03-23 11:03:33	CSV Re-run Delete
Files Before Dec 1 2020	Success	2021-03-09 12:52:16	CSV Re-run Delete
Files >= 500MB	Success	2021-03-04 14:25:37	CSV Re-run Delete
Files older than 12months	Success	2021-03-02 13:15:00	CSV Re-run Delete
500 Oldest Files	Success	2021-02-23 17:44:18	CSV Re-run Delete
2019 Files	Success	2021-02-17 13:51:46	CSV Re-run Delete
Files accessed in last 30 days	Success	2021-02-10 10:28:35	CSV Re-run Delete

Page indicator: 1 - 7 of 7

Figure 40: Saved Reports

Smart Tier

Administrators with large scale Network Attached Storage (NAS) environments frequently look for ways to improve storage efficiency, simplify administration, and decrease cost. With Nutanix Files 4.0, we added a native tiering framework called Smart Tier to address these requirements. Smart Tier enables you to move cold or rarely used data to lower-cost storage while maintaining a single namespace. You can tier data to any qualified S3 API compliant target, including on-premises solutions or the public cloud. Smart Tier enables hybrid multicloud environments to handle unstructured data to retain long term, save money, or provide a virtually limitless pool of local storage for your Nutanix Files workloads.

Smart Tier Architecture

The Nutanix Files 4.0 tiering engine uses APIs over HTTPS to accept tiering and recall requests and to move data to S3-compliant targets. Three targets have been validated for the initial release: Nutanix Objects, Amazon S3 Standard and IA tiers, and Wasabi Cloud Storage. When you tier files, Nutanix Files maintains the metadata for the file while it moves the data to the S3 target.

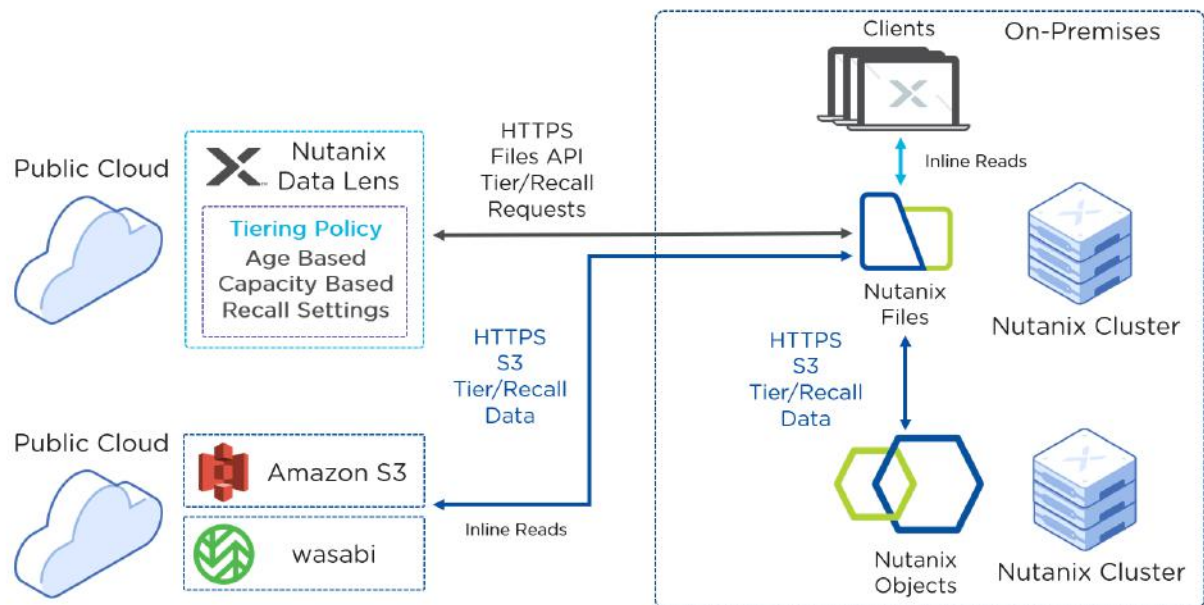


Figure 41: Smart Tier Architecture

Smart Tier maintains a file stub in the SMB or NFS share path where clients can perform inline reads to access the data. You can also recall files automatically, based on access patterns, or with manual recall operations. Configure tiering policies through Nutanix Data Lens, which allows you to set age- and capacity-based thresholds, manual or automatic tiering schedules, and recall settings.

Nutanix Data Lens

Data Lens is a Nutanix software as a service (SaaS) that provides file analytics and reporting, anomaly detection, audit trails, ransomware protection features, and tiering management for your entire Nutanix Files environment.

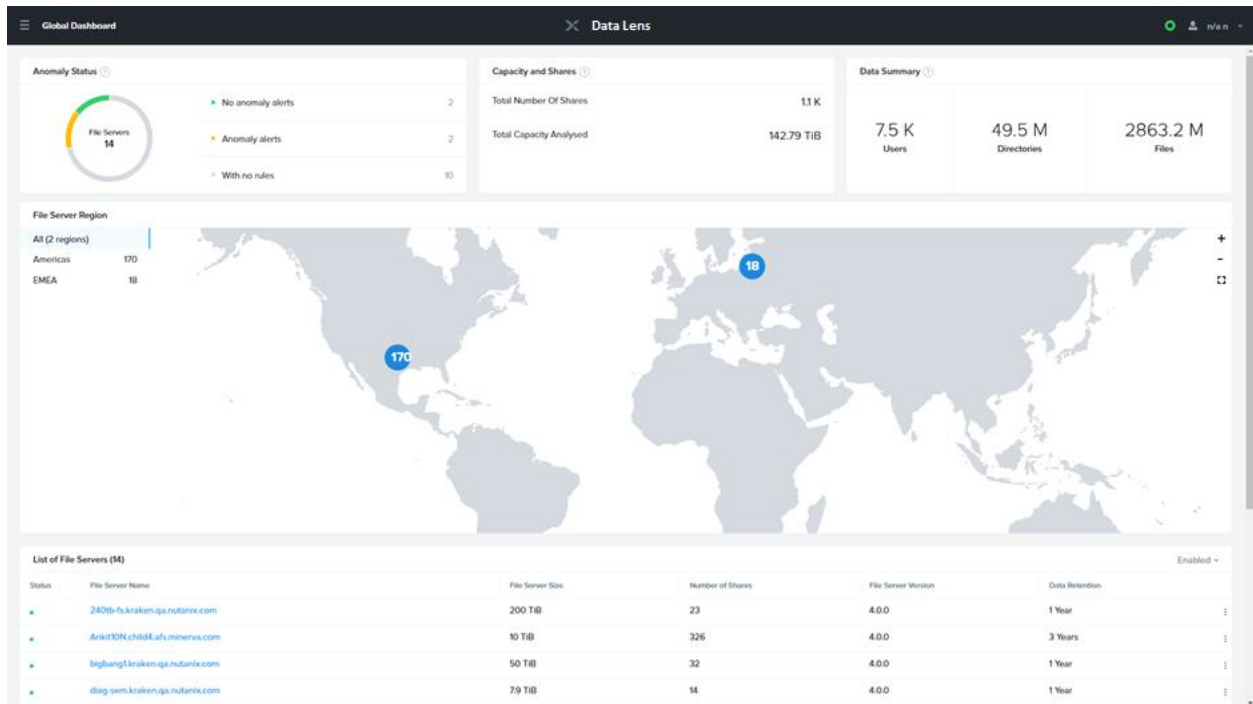


Figure 42: Data Lens Dashboard

Data Lens tracks the access patterns for all shares and exports to determine the age of each file for a given file server instance. You can define hot, warm, and cold data age categories to match your requirements. The data age explorer lets you view the file age on a share-by-share basis to understand what data is included in a given tiering policy.

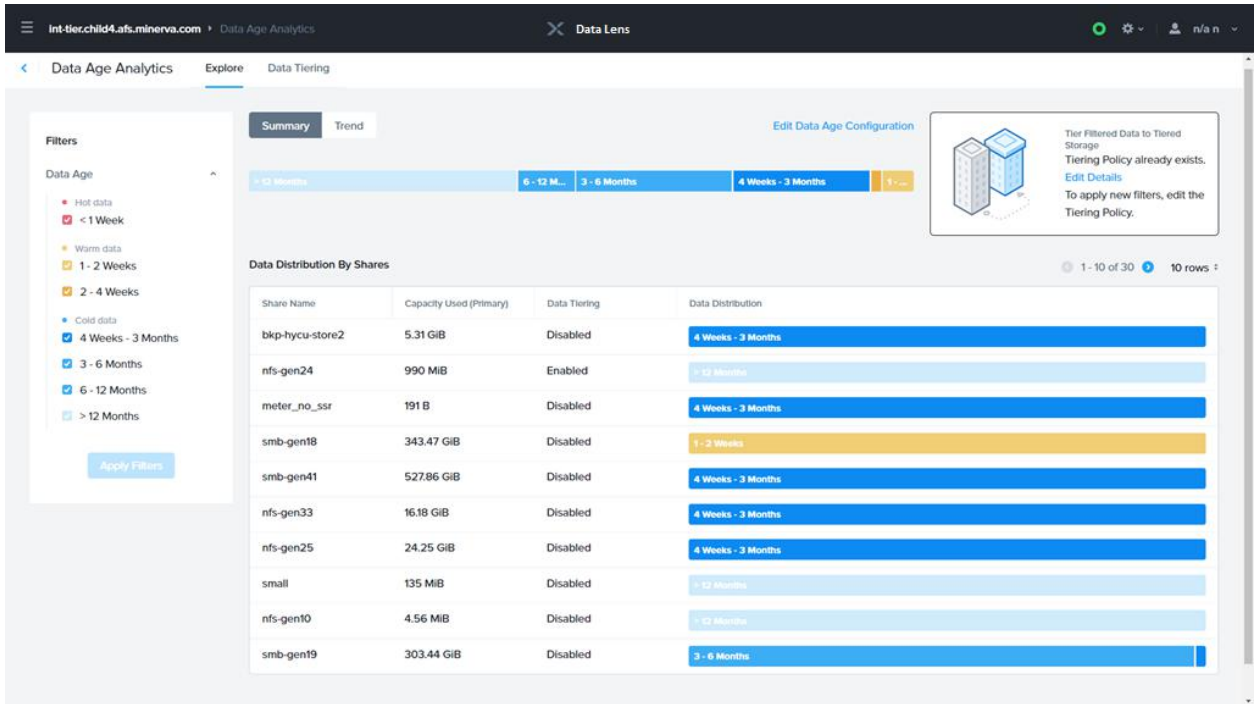


Figure 43: Data Lens Age Explorer

Smart Tier Configuration

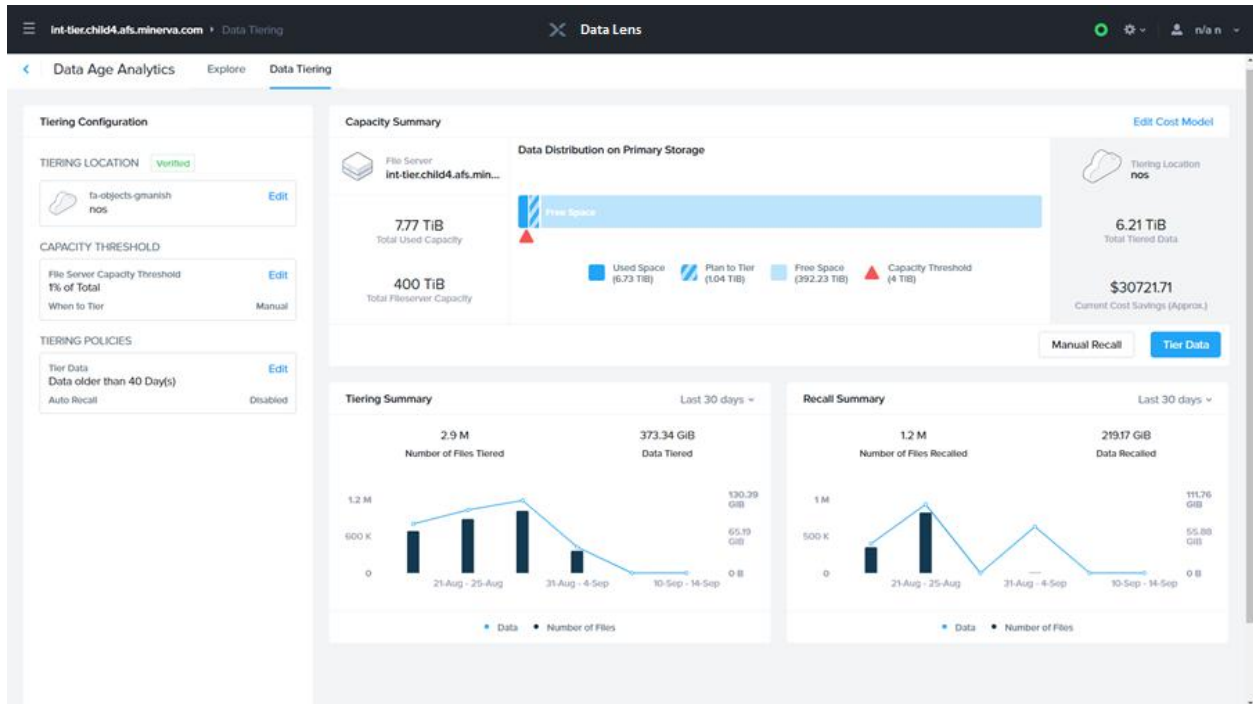


Figure 44: Data Lens Tiersing

To configure Smart Tier, define a tiering location (a supported S3-compliant target). Configure the following:

- Target URL over HTTPS
- Bucket name
- Access and secret keys
- Retention period
- Certificate (optional)

The retention period determines when Nutanix Files deletes a tiered object after removing the stub from the file system.

After you configure your location, define a capacity threshold and when-to-tier policy. The capacity threshold represents the percent of allocated space consumed before the system considers tiering. You also define whether the

system performs the tiering policy manually or on an automated basis, and the time windows in which automated tiering occurs.

The third step is to define the formal tiering policy, including:

- Age of the data (based on the last file read or write operation)
- Minimum file size (64 KB or larger)
- Shares to exclude from the policy
- Automatic recall settings

The automatic recall setting lets you define whether to recall files automatically and under what conditions, specifically the number of times the file is accessed over a given period. If you choose to recall files manually, Data Lens lets you search for and choose individual files, folder paths, or shares to recall.

Tiered Data

When data reaches the configured age and the file server is at its capacity threshold, Smart Tier moves file data to the target. Smart Tier marks the previously consumed space as free, creating additional space in the share for new or recalled data.

Tiered files appear as regular files in their folder paths. SMB shares have an offline attribute set, which in Windows Explorer shows an X beside the file icon.




Figure 45: Tiered File

You can see the file size (size on disk), which represents just the metadata, and the actual file size, which represents the tiered data. You can read (which here means retrieve the data inline from) tiered files at any time, but you can't write

to a tiered file directly, so you need to either retrieve the file or make a separate copy to edit it.

Through Data Lens you can see the hot, warm, cold, and archived data associated with tiering. You can also define cost thresholds that help you see the current and potential cost savings based on your tiering policies.

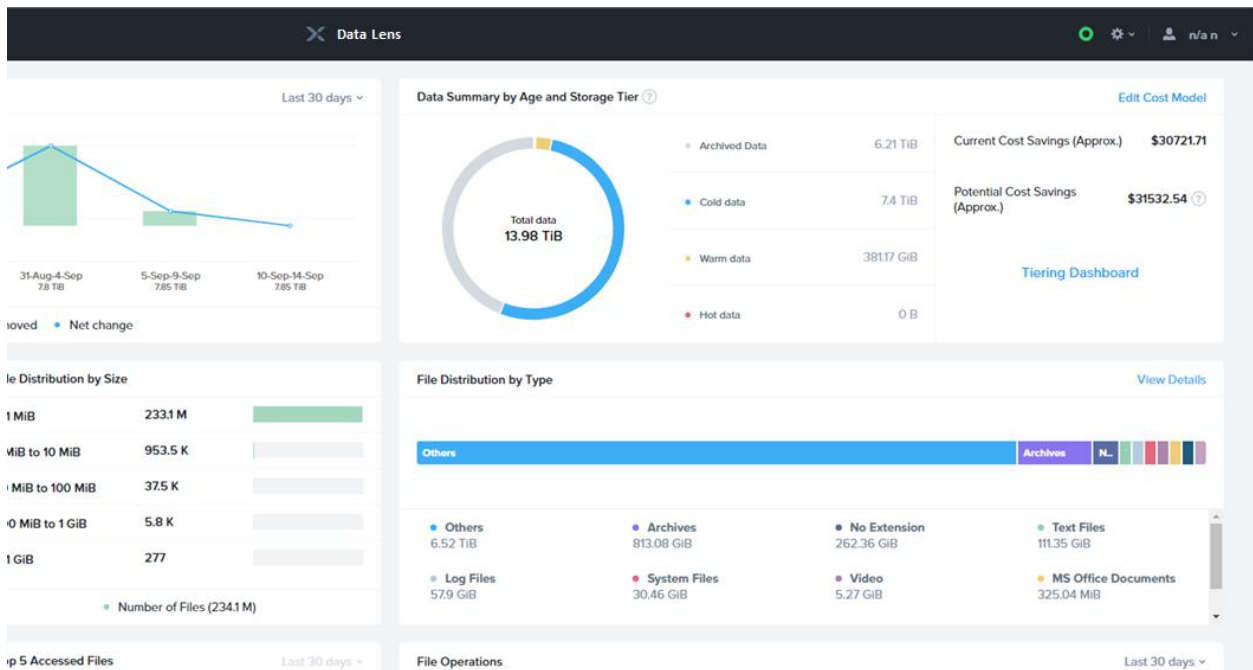


Figure 46: Data Summary by Age and Storage Tier

Hypervisor-Specific Support

Nutanix supports ESXi and AHV for Files. For ESXi support, you need vCenter credentials to deploy Nutanix Files and to create DRS rules to make sure the FSVMs are on different nodes. You must register vCenter with the Nutanix cluster instances where you deployed Files. The deployment process generates the DRS rules for AHV automatically.

Data-at-Rest Encryption

Nutanix Files supports data-at-rest encryption (DARE) with self-encrypting drives (SEDs) or software-defined DARE using AOS. Official support for AOS

software encryption with Nutanix Files begins with AOS 5.10.2 in combination with Nutanix Files 3.5.

For AHV, you enable software encryption at the cluster level. For ESXi, you enable software encryption at either the cluster level or the storage container level. To enable storage container-level software encryption with Nutanix Files you must use nCLI:

```
<ncli> storage-container edit enable-software-encryption=true  
name=<Files_Container_Name> force=true
```


5. Backup and Disaster Recovery

Following modern data protection methodologies, Nutanix provides administrators and users quick restore access using Self-Service Restore (SSR) and site recovery with Nutanix-based snapshots.

Self-Service Restore

Administrators can enable SSR at any time for SMB or NFS shares. Windows Previous Version in each folder exposes SSR for SMB shares. SSR for NFS shares is exposed as a hidden snapshot directory in each folder.

Protection configuration: filesvr1 ? x









Self Service Restore

Snapshots will be created based on the configured schedule for shares/exports with self service restore enabled. End users can access these snapshots through their native interface.

Enabled for

Self service restore is enabled for all shares/exports. ✓

1 hour RPO • Total :38 Snapshots + Add schedule

TYPE	FREQUENCY	SNAPSHOTS	ACTIONS
Hourly	Every 1 hour	24	 
Daily	Every 1 day	7	 
Weekly	Every week on Sun	4	 
Monthly	Every month on 1	3	 

Note: Self Service Restore can be disabled for specific shares/exports if required

Figure 47: Protection Using SSR

SSR allows you to view snapshots of shares while the share is in use. The share snapshots are read-only, point-in-time copies.

You can view and restore removed or overwritten files, which allows you to choose a share snapshot from the same file at different times during the file's history. Administrators can configure a snapshot schedule at the file-server level that applies to all shares in the file server. Nutanix Files supports 24-hour, daily, weekly, and monthly snapshots on a fixed schedule. The default snapshot policy includes:

- Hourly (24 per day).
- Once daily for seven days.
- Four weekly.
- Three monthly.

Note: You can configure a maximum of 50 total snapshots in the SSR schedule.

You can change schedule frequency to suit your requirements, including shorter intervals for same-day protection against accidental deletions. You can enable SSR during or after share creation. After share creation, the administrator can change the current settings using the share update workflow feature. SSR supports both standard and distributed SMB shares.

Protection Domains and Consistency Groups

Nutanix provides integrated, automated disaster recovery between Nutanix clusters. A Nutanix Files cluster can be protected with Prism and uses the same asynchronous replication with protection domains and consistency groups as any other Nutanix cluster. A protection domain is a defined group of entities (VMs and volume groups) that you back up locally on a cluster and that may replicate to one or more remote sites. A consistency group is a subset of the entities in a protection domain. Consistency groups are configured to snapshot a group of VMs or volume groups in a crash-consistent manner.

Note: Nutanix Files supports both asynchronous (RPO of one hour or longer) and NearSync (down to one minute RPO) schedules. NearSync support requires a minimum of AOS 5.11.1 and Files 3.6. Use AOS node sizing guidelines when you determine your RPO requirements.

When you create a file server, Prism automatically sets up a corresponding protection domain, which it annotates with the Nutanix Files cluster name.

Prism also creates multiple consistency groups in a protection domain, including a group that includes all FSVMs.

Once you've protected Nutanix Files, all future operations on it (such as adding or removing FSVMs or adding or deleting volume groups) automatically update the corresponding consistency group in the protection domain.

Cluster Migration, Failure, and Restoration

In the event of a Nutanix Files cluster failure, restore Files in a remote cluster by initiating the Activate workflow, which restores from the last good snapshot. If you're moving your file services because you need to shut the cluster down, as with a planned outage, the Migrate workflow shuts down all the FSVMs and takes a final snapshot for replication. You run the Migrate workflow from the Nutanix cluster that owns the active protection domain. You initiate both the Activate and Migrate workflows from the Data Protection menu in Prism.

After you run either the Activate or Migrate workflow, you also need to activate the file server instance, which you can do from the File Server menu in Prism. Activating the file server may require you to configure network VLANs on the replica site before Nutanix Files becomes operational again. These VLANs can be in subnets that are the same as or different from the source site. The recovery process is like creating the file server but in this context you can change networks and IP addresses if necessary.

Cloning

Because Nutanix Files cloning doesn't affect the original Files cluster, it offers improved support for a variety of use cases:

- Backups at the primary and secondary sites.
- Disaster recovery at the secondary site.
- File server recovery from a specific point in time.
- File server creation at the primary or remote site for testing or development.
- File server clone copies.

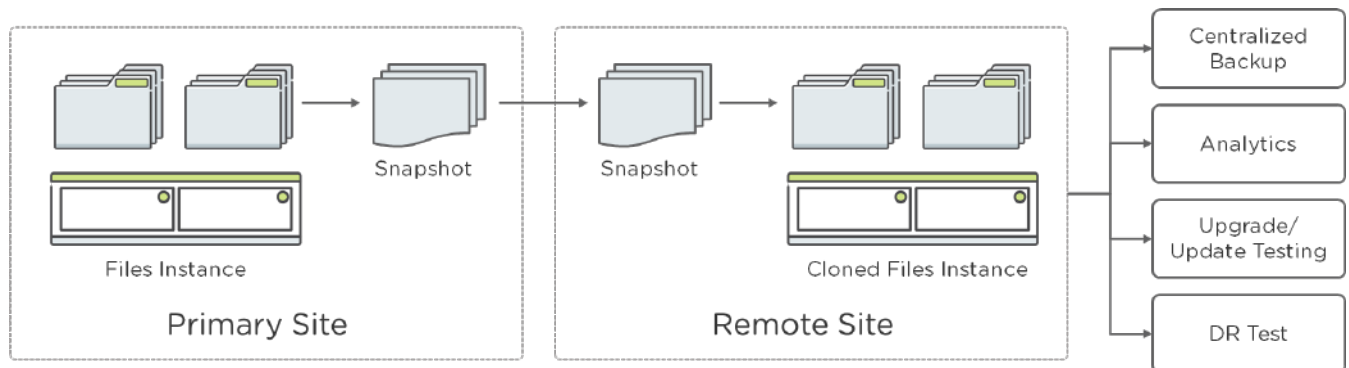


Figure 48: Nutanix Files Cloning Use Cases

Files uses Nutanix native snapshots to clone entire file servers. The clone is a thin copy that consumes minimal storage space. File server clones reside on the same container as the original and maintain the original security permissions. During the clone process you can specify new IP addresses and give the cloned file server a new name.

Files Smart Disaster Recovery

Nutanix Files long relied on the core snapshot and remote replication capabilities of Nutanix AOS software. Using the core Nutanix software had several benefits, including simplified and consolidated administration that aligned with the core hyperconverged infrastructure (HCI) environment. But there were also some drawbacks around granularity, node density, and the active-passive nature of share access and failover orchestration. Files Smart DR helps address these challenges while maintaining the same simple and consolidated administrative experience.

Prism Central Integration

Files Smart DR moves remote replication management and orchestration into Prism Central. With recent releases of Prism Central, Nutanix Files has supported an integrated service called Files Manager. The first release of Files Manager discovers all Files instances running on clusters registered in Prism Central and provides views on all alerts and events across your file server farm. You can also view the file server configurations and launch Prism Element to manage the file servers.

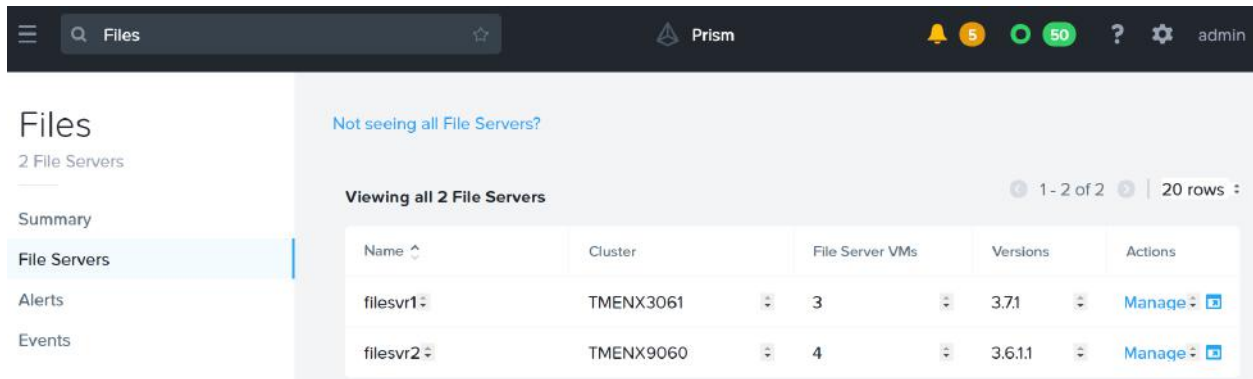


Figure 49: File Manager in Prism Central

Files Manager 2.0 provides a data protection menu where you can configure, monitor, and orchestrate failover and failback operations for Smart DR.

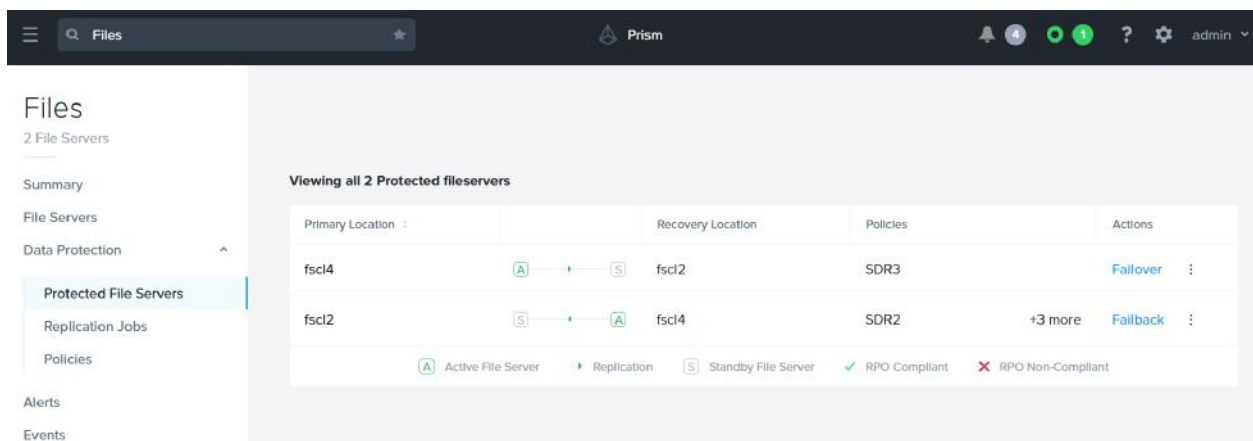


Figure 50: Files Manager Data Protection

Files Smart DR Architecture

In addition to moving management into Prism Central, Smart DR changes several key areas of Nutanix Files remote replication. First, the AOS clusters no longer manage the replication engine itself using native Nutanix protection domains; instead, Nutanix Files manages replication directly. As with Files Self-Service Restore (SSR), Smart DR takes snapshots at the share (file system) level, replicating block-level incremental changes between source and target.

Replication occurs between active file servers running on their respective Nutanix clusters. Shares that function as replication targets are available in a

read-only state. Replicating between active file servers helps shorten failover times, reduce your recovery time objective (RTO), and simplify use of the replicated data for use cases like backup consolidation or reporting.

Because Files now manages replication, node density limits specific to AOS snapshots and replication no longer apply. You can now use our most storage-dense nodes, which support up to 350 TB of hybrid storage, with the benefits of native remote replication.

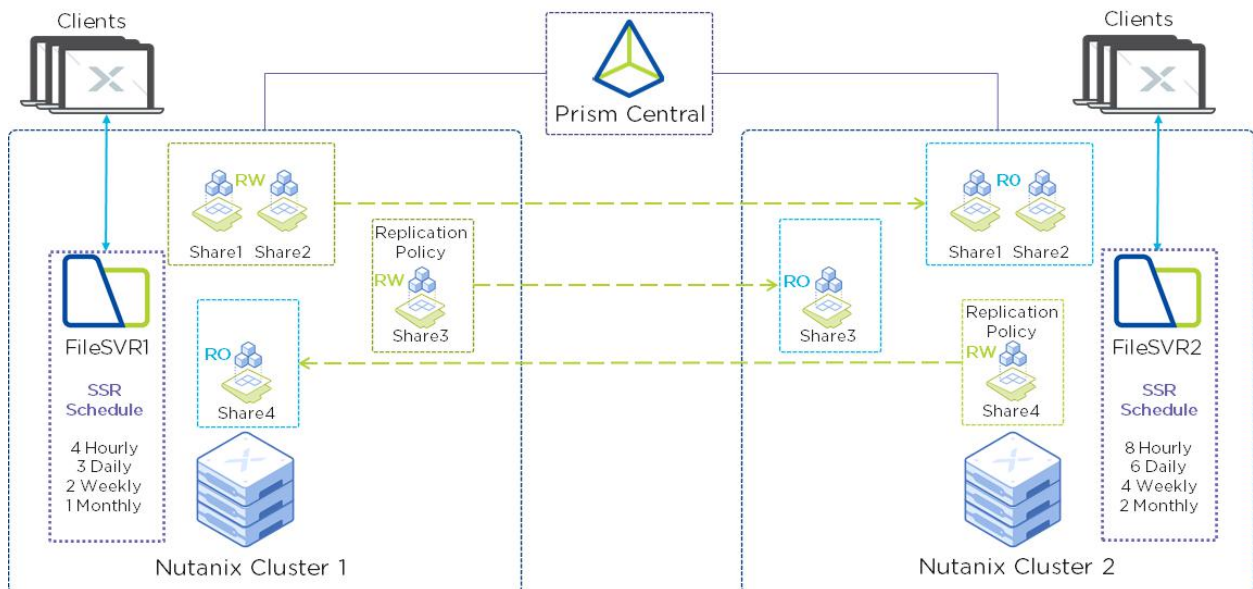


Figure 51: Files Smart DR

Further, for the file systems that support the shares performing replication, you can now set policies on a share-by-share basis to manage your recovery point objective (RPO) at the share level instead of at the file server level.

Replication Policies

A replication policy defines the share or group of shares you want to replicate. The policy also defines the source and target file servers and the replication frequency. You can create multiple policies as needed and specify a default policy for any newly created shares.

Create a Data Protection Policy

1 Primary and Secondary Location 2 Settings

Primary Location (Source File Server)

Select File Server
fsc12

Shares to be protected
1 of 7 shares. [Edit](#)

Replication Schedule

Recovery Point Objective (RPO)
10 minutes

A snapshot of the share will be taken and replicated to the target every 10 minutes

Start immediately
 Start from a specific point in time

Recovery Location (Target File Server)

For every protected share, a share will be created on the target with read only access

Select File Server
fsc14

Figure 52: Data Protection Policy

You can configure your replication schedule to be as short as one minute.

Monitoring

Files Manager in Prism Central provides SLA monitoring and job replication status at several levels. You can use the RPO compliance overview provided on the summary page to quickly understand if replication is keeping up with the defined policy.

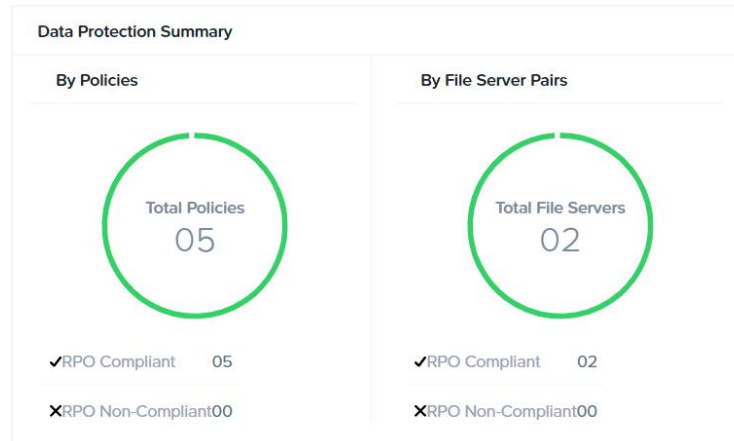


Figure 53: Data Protection Compliance

You can also view each replication job to monitor its completion percentage, start and end times, amount of data synchronized, and average network bandwidth usage.

Failover and Failback

You can manage planned or unplanned failover and failback from Prism Central. During failover operations, Prism Central orchestrates the required updates to the DNS and Active Directory service principal names (SPNs) to move the file server instance name from the source to the target.

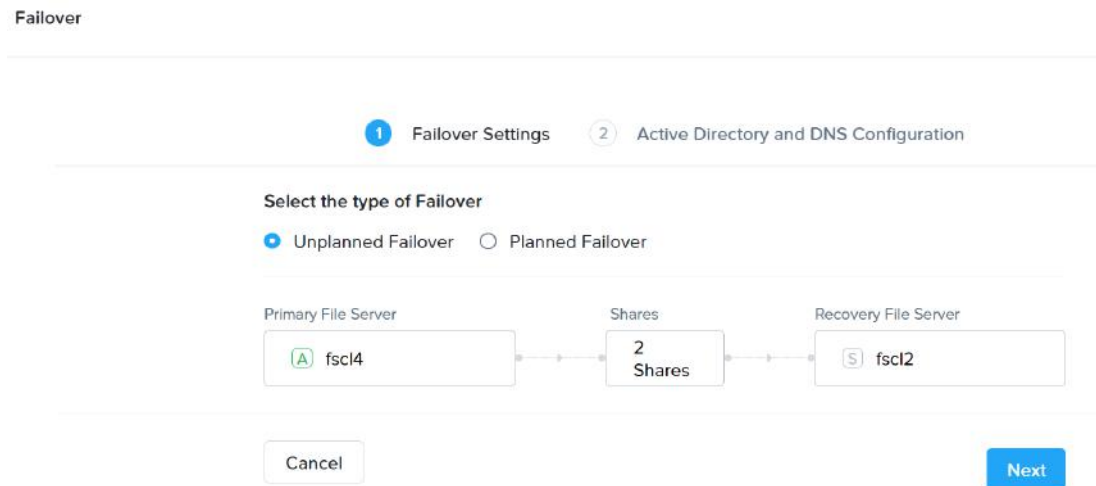


Figure 54: Files Smart DR Failover

With a planned failover, you can choose to begin replicating in the opposite direction automatically. Replicating after failover helps maintain service level agreements (SLAs) and RPOs during your failover testing or disaster avoidance operations.

SSR Interoperability

With Files 4.0, Smart DR supports replicating the snapshots between the source share and its target. The target share retention schedule doesn't have to match the source retention schedule. For example, if the file server is configured on the source to maintain the last two hourly snapshots, you can configure the target file server to maintain the last ten hourly snapshots to provide a longer retention window. Alternatively, you can match retention schedules to ensure that the same SSR copies exist in both sites for any failover event.

SSR schedules must be aligned by their frequency type between the source and target. For example, if you have daily snapshots scheduled on your source, you need daily snapshots scheduled on your target if you want them to be replicated and retained.

Files Smart DR Summary

Files Smart DR is a smart, simple, and effective way to replicate between Files instances, either on-premises or running on Nutanix Clusters in AWS (NCA).

In summary, Files Smart DR gives you the following advantages—all while maintaining simple and streamlined administration for your Nutanix environment:

- Prism Central integration.
- Share-level replication policies.
- Replication to an active read-only file server target.
- Faster RTOs.
- Support for storage-dense nodes.
- SSR compatibility.

6. Third-Party Integration

Antivirus

To protect users from malware and viruses, you need to address both the client and the file server. Nutanix currently supports third-party vendors that use Internet Content Adaptation Protocol (ICAP) servers. ICAP, which is supported by a wide range of security vendors and products, is a standard protocol that allows you to integrate file and web servers with security products. Nutanix chose this method to give customers wide latitude in selecting the antivirus solution that works best for their specific environments.

Following is the workflow for an ICAP-supported antivirus solution:

1. An SMB client submits a request to open or close a file.
2. The file server determines if the file needs to be scanned, based on the metadata and virus scan policies. If a scan is needed, the file server sends the file to the ICAP server and issues a scan request.
3. The ICAP server scans the file and reports the scan results back to the file server.
4. The file server takes an action based on the scan results:
 - a. If the file is infected, the file server quarantines it and returns an access denied message to the SMB client.
 - b. If the file isn't infected, it returns the file handle to the SMB client.




Figure 55: ICAP Workflow

The ICAP service runs on each Nutanix Files file server and can interact with more than one ICAP server in parallel to support horizontal scale-out of the antivirus server. The scale-out nature of Files and one-click optimization greatly mitigates any antivirus scanning performance overhead. If the scanning affects Nutanix Files FSVM performance, one-click optimization recommends increasing the virtual CPU resources or scaling out the FSVMs. This feature also allows both the ICAP server and Files to scale out, ensuring fast responses from the customer's antivirus vendor.

Tip: We recommend configuring two or more ICAP servers for production.

Antivirus Setup SennaFS ? X

ICAP Servers · Scan Settings

Connect ICAP servers for antivirus integration. Two or more ICAP servers recommended.

LIST OF ICAP SERVERS [+ Connect ICAP Server](#)

IP ADDRESS OR HOSTNAME	PORT	DESCRIPTION	CONNECTION STATUS	ACTIONS
10.7.69.224	1344		✓ OK	✎ · ✕
10.7.69.225	1344		✓ OK	✎ · ✕
10.5.236.201	1344		✓ OK	✎ · ✕
10.7.69.226	1344		✓ OK	✎ · ✕

[Next](#)

Figure 56: Configure Multiple ICAP Servers

Nutanix Files sets scanning defaults across the entire file server. You can enable scan on write and scan on read. Scan on write begins when the file is closed, and scan on read occurs when the file is opened. You can also exclude certain file types and files over a certain size. Share scan policies can override any defaults set for the file server.

The screenshot shows the 'Antivirus Setup SennaFS' window with the 'Scan Settings' tab selected. The settings are as follows:

- Scan Settings (For all shares)**: Settings can be overridden for individual share if required, through the share level antivirus settings.
- On access scan type**:
 - SCAN ON WRITE
 - SCAN ON READ
- EXCLUDE FILE TYPES**: Comma separated extension like .db, .txt to be excluded from scanning
- EXCLUDE FILES LARGER THAN**: No file size limit (MiB)
- SHOW ADVANCED OPTIONS
- SCAN TIMEOUT**: 60 seconds (MAXIMUM 240 SECONDS)
- BLOCK ACCESS TO FILES IF ANTIVIRUS SCAN CANNOT BE COMPLETED (RECOMMENDED)

Buttons: Back, Cancel, Save

Figure 57: Default Scan Settings

For each ICAP server, Nutanix Files spins up no more than 10 parallel connections per FSVM and randomly distributes the file scanning among all the ICAP servers. With heavier workloads, which may involve many scan requests and use all connections, you can give the scan servers more processing power to scan more files. As soon as the current scan finishes, the scan server picks up the next file from the queue, which keeps the number of active connections at 10.

ICAP Servers	Reports	Quarantined Files	Unquarantined Files ⓘ																																												
Actions ▾ 1 of 17 files selected. 1-10 of 17 < > ⚙ search in table 🔍																																															
Rescan Unquarantine Delete	<input type="checkbox"/> General_1 <input type="checkbox"/> roaming <input type="checkbox"/> General_1 <input type="checkbox"/> General_38 <input type="checkbox"/> roaming <input type="checkbox"/> General_38 <input type="checkbox"/> roaming <input type="checkbox"/> roaming	<table border="1"> <thead> <tr> <th>FILE PATH</th> <th>THREAT DESCRIPTION</th> <th>ICAP SERVER</th> <th>SCAN TIME</th> </tr> </thead> <tbody> <tr> <td>/vdi2.V2/Documents/virus - Copy (4).txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.226</td> <td>08/03/17, 2:03:00 PM</td> </tr> <tr> <td>/virus1.txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.224</td> <td>08/02/17, 1:52:34 PM</td> </tr> <tr> <td>/yifeng_virus2</td> <td>EICAR-Test-File;</td> <td>10.5.236.201</td> <td>08/07/17, 3:35:25 PM</td> </tr> <tr> <td>/vdi3.V2/Music/roaming_virus - Copy (3).txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.225</td> <td>08/03/17, 2:14:00 PM</td> </tr> <tr> <td>/yifeng_virus3.txt</td> <td>EICAR-Test-File;</td> <td>10.5.236.201</td> <td>08/07/17, 3:36:37 PM</td> </tr> <tr> <td>/virus9.txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.226</td> <td>08/09/17, 10:14:56 AM</td> </tr> <tr> <td>/vdi2.V2/Documents/virus - Copy (2).txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.226</td> <td>08/03/17, 2:03:00 PM</td> </tr> <tr> <td>/vsample04.txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.225</td> <td>08/07/17, 5:52:28 PM</td> </tr> <tr> <td>/vdi39.V2/Documents/virus.txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.226</td> <td>08/04/17, 7:02:19 PM</td> </tr> <tr> <td>/vdi3.V2/Music/roaming_virus - Copy (4).txt</td> <td>EICAR-Test-File;</td> <td>10.7.69.226</td> <td>08/03/17, 2:14:00 PM</td> </tr> </tbody> </table>	FILE PATH	THREAT DESCRIPTION	ICAP SERVER	SCAN TIME	/vdi2.V2/Documents/virus - Copy (4).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:03:00 PM	/virus1.txt	EICAR-Test-File;	10.7.69.224	08/02/17, 1:52:34 PM	/yifeng_virus2	EICAR-Test-File;	10.5.236.201	08/07/17, 3:35:25 PM	/vdi3.V2/Music/roaming_virus - Copy (3).txt	EICAR-Test-File;	10.7.69.225	08/03/17, 2:14:00 PM	/yifeng_virus3.txt	EICAR-Test-File;	10.5.236.201	08/07/17, 3:36:37 PM	/virus9.txt	EICAR-Test-File;	10.7.69.226	08/09/17, 10:14:56 AM	/vdi2.V2/Documents/virus - Copy (2).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:03:00 PM	/vsample04.txt	EICAR-Test-File;	10.7.69.225	08/07/17, 5:52:28 PM	/vdi39.V2/Documents/virus.txt	EICAR-Test-File;	10.7.69.226	08/04/17, 7:02:19 PM	/vdi3.V2/Music/roaming_virus - Copy (4).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:14:00 PM	
FILE PATH	THREAT DESCRIPTION	ICAP SERVER	SCAN TIME																																												
/vdi2.V2/Documents/virus - Copy (4).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:03:00 PM																																												
/virus1.txt	EICAR-Test-File;	10.7.69.224	08/02/17, 1:52:34 PM																																												
/yifeng_virus2	EICAR-Test-File;	10.5.236.201	08/07/17, 3:35:25 PM																																												
/vdi3.V2/Music/roaming_virus - Copy (3).txt	EICAR-Test-File;	10.7.69.225	08/03/17, 2:14:00 PM																																												
/yifeng_virus3.txt	EICAR-Test-File;	10.5.236.201	08/07/17, 3:36:37 PM																																												
/virus9.txt	EICAR-Test-File;	10.7.69.226	08/09/17, 10:14:56 AM																																												
/vdi2.V2/Documents/virus - Copy (2).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:03:00 PM																																												
/vsample04.txt	EICAR-Test-File;	10.7.69.225	08/07/17, 5:52:28 PM																																												
/vdi39.V2/Documents/virus.txt	EICAR-Test-File;	10.7.69.226	08/04/17, 7:02:19 PM																																												
/vdi3.V2/Music/roaming_virus - Copy (4).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:14:00 PM																																												

Figure 58: Quarantined Files

Once Nutanix Files quarantines a file, the administrator can rescan the file, remove it from quarantine, or delete it. You can search quarantined files if you need to restore a file quickly.

If your antivirus vendor doesn't support ICAP, you can scan the shares by installing an antivirus agent on a Windows machine, then mounting all the shares from the file server. This approach allows you to schedule scans for periods of low usage. At the desktop or client level, you can set your antivirus solution to scan on write or scan only when files are modified. You can configure high-security environments to scan inline for both reads and writes.

See the [Compatibility and Interoperability Matrix](#) on the Nutanix portal for the latest list of qualified ICAP server vendors with Nutanix Files.

File Operations Monitoring

File operations monitoring is a native Nutanix Files API used to forward all SMB and NFS operations run by clients to a registered target repository.

File operations monitoring can be broken down into two major areas for third-party vendors:

1. File activity

2. Audit

Partner software can register with the Nutanix Files instance to capture file activity events. Partner software can also make REST calls to Files to create a policy defining which file notification types you want to receive and the shares you want to receive them from. The partner software uses a web client and communicates with the Nutanix Files file server using HTTPS requests. The HTTPS communication relies on SSL authentication. The HTTP server runs with either its unique self-signed SSL certificate or a Transport Layer Security (TLS) connection with the partner, then exchanges the keys between the two parties and sends messages to the partner server over this secure channel.

Third-party software can use different protocol communication methods such as syslog, Google, or Kafka protobuf. One common way to use the Nutanix file operations monitoring API is to forward events to a syslog server for any future auditing needs. Several vendors have integrated with this operations monitoring API:

- Peer Software

Peer Software has integrated with Nutanix, including the file operations monitoring API, to support active-active deployments with real-time replication and file locking between heterogenous environments. To learn more about the Peer Software and Nutanix integration, see the [PeerGFS and Nutanix Files datasheet](#).

- Netwrix

Netwrix has integrated with the Nutanix file operations monitoring API. Netwrix provides comprehensive visibility into changes and data access across Nutanix Files instances. Netwrix offers an add-on for Nutanix Files. See [netwrix.com](https://www.netwrix.com) for more details.

- Varonis

Varonis version 8.6 introduces support for Nutanix Files version 3.7.1 and higher. This support includes Varonis's entire platform, DatAdvantage, DatAlert/DatAlert Analytics, DataPrivilege, Data Classification Engine, DatAnswers, Automation Engine, and Data Transport Engine.

Intelligent Backup

Traditional NAS backup approaches like Network Data Management Protocol (NDMP) have core limitations. Problems like complexity, performance (including the need for periodic full backups), and scale led Nutanix to develop a more modern backup solution: a changed file tracking (CFT) API that third-party backup vendors use. This API shortens backup times because it doesn't perform a metadata scan across your file server, which could contain millions of files and directories. The API enables scale as backups can occur across multiple FSVMs in parallel. The API also ensures that customers aren't locked into any one solution as backups taken using the Nutanix Files API can be restored to other vendors.

CFT uses internal snapshots to track differences between backup windows. The backup software vendor requests the creation of these snapshots, and Files returns a list of the changed files in the form of a share. The backup software proceeds to mount the share onto the backup server to be read for backup.

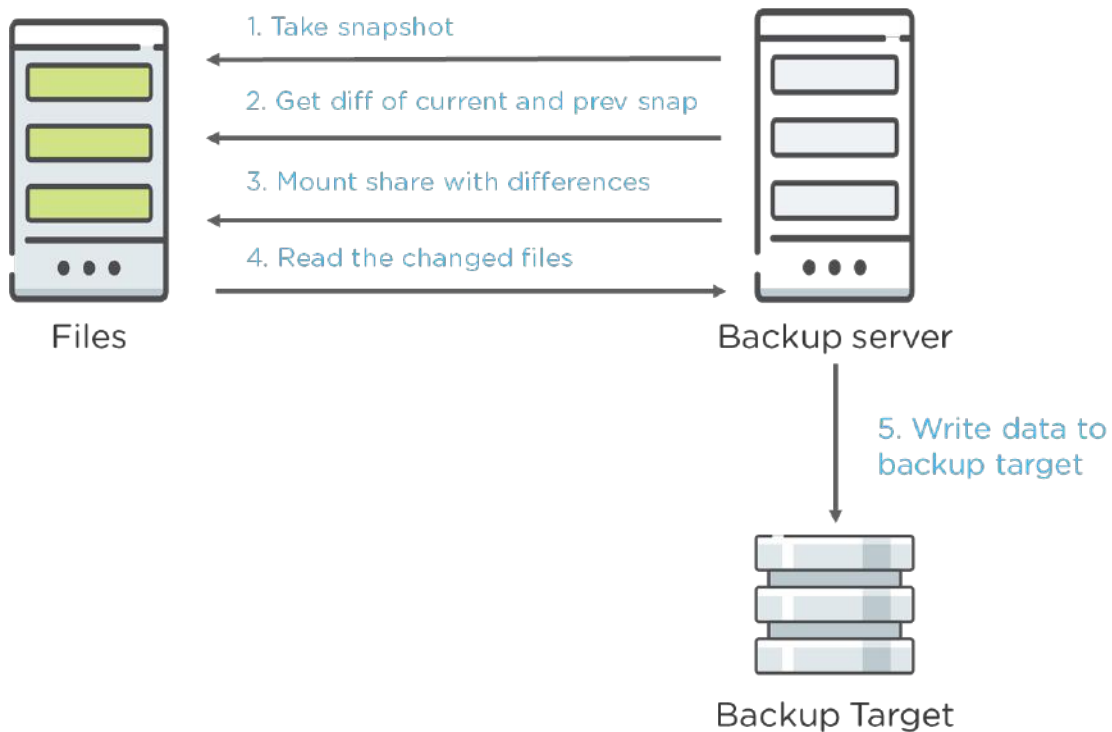


Figure 59: CFT Backup Process

Backup software can specify multiple shares and their respective snapshot information. Nutanix Files returns the list of URLs that map to the number of client streams that can start in parallel. Because shares are distributed evenly across the FSVMs based on load, the backup software can take advantage of all the FSVMs to drive throughput.

HYCU Backup and Recovery is integrated with the Nutanix Files CFT API. Commvault version 11 service pack 15 is integrated with the Nutanix Files CFT API as of the 3.5 release. Veritas Netbackup 9.1 is fully integrated with the Nutanix Files CFT API.

There are two ways to back up Files shares with software that doesn't support CFT. One option is to run the backup application on a Windows machine and map the UNC path of the share as a drive that needs to be backed up.

There are also vendors that provide support for backing up file shares without mounting to a guest VM. These applications can read directly from the UNC path.

Both Veeam and Rubrik are validated as non-CFT backup solutions for Nutanix Files. Refer to the [Compatibility and Interoperability Matrix](#) for the full list of validated backup applications.

Tip: Don't try to back up the FSVMs. FSVMs are stateless and data is stored in Nutanix volume groups. Back up Nutanix Files from the share level. You can protect FSVMs using Nutanix protection domains.

Because the system spreads different standard shares across the cluster, try to back up multiple shares at the same time with multiple subclients. The distributed share allows you to configure multiple data readers to drive throughput.

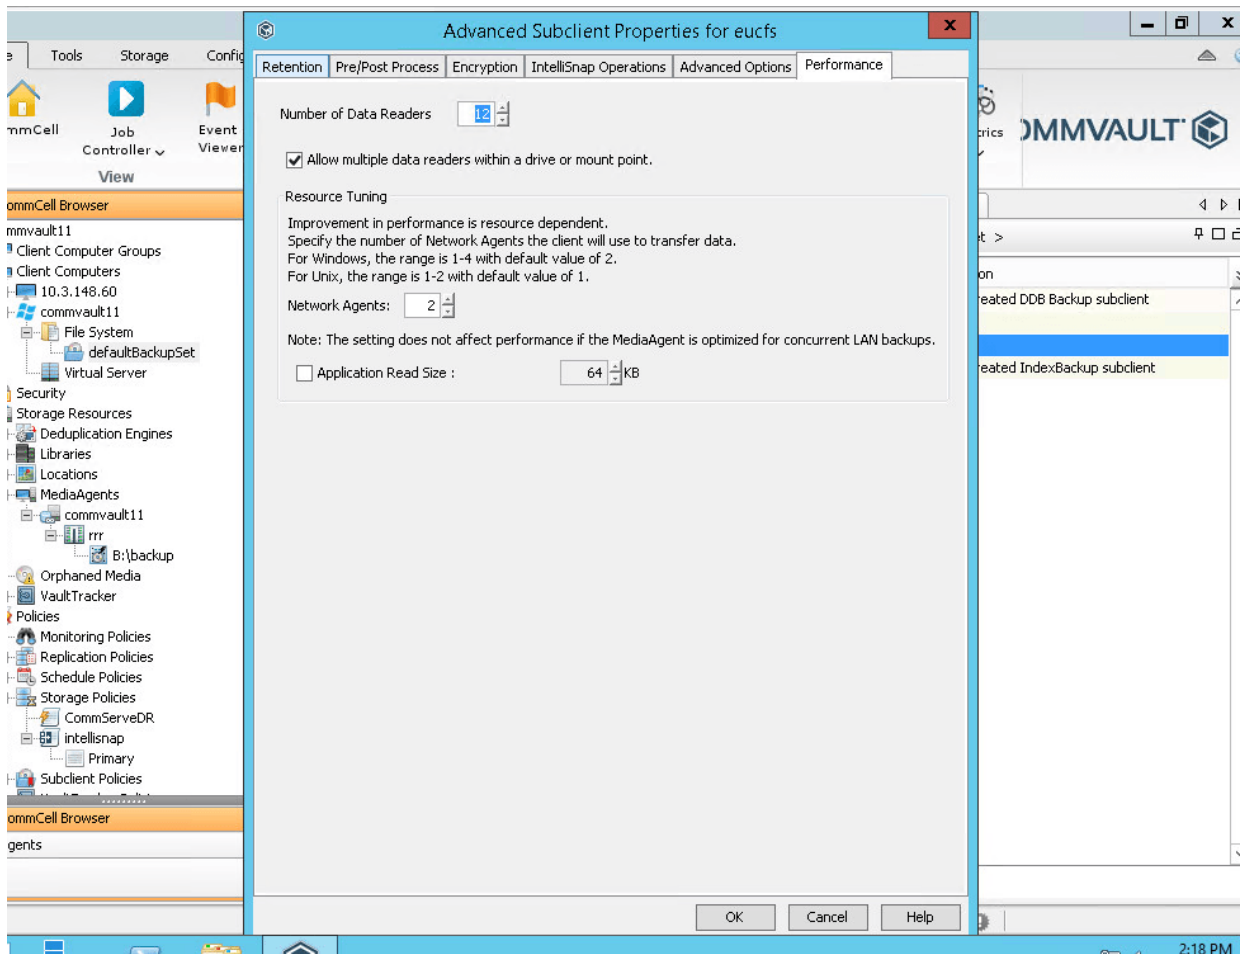


Figure 60: Adding More Readers to Drive Throughput on a Home Share

As an example of backup using Commvault, we tested 400 users spread out on three FSVMs, placing the data on a home share. We found that adding more readers for the backup job could increase performance. The bottleneck was the media agent, which was the backup destination for the files. The media agent was virtualized and configured with only eight vCPU. Adding more vCPUs to the media agent achieved a shorter backup time.




Figure 61: Backup Job Details

As the previous figure shows, we achieved almost 600 GB average throughput per hour with a fairly small setup.

7. Conclusion

Nutanix Files delivers on-demand performance and automated provisioning to provide highly scalable file management. It reduces the administration and configuration time needed to deploy and maintain your environment, providing a public cloud experience within your private cloud.

AOS distributed storage and Prism make hosted file services highly resilient and easy to use—for example, you can configure the native data efficiency features of AOS, such as erasure coding (EC-X), for each individual file server. Prism lets you administer network and resource management, Active Directory, fault tolerance, and Nutanix Files share and export management all in one place, vastly improving operational efficiency.

File Analytics gives you deep insight into your data, including storage consumption trends, file sizes, file counts, and data age. Analytics also provides end-to-end audit trails and anomaly detection. Ransomware intelligence helps you detect, protect against, analyze, and recover from ransomware threats.

Because you can deploy Files on a new or existing Nutanix environment, as well as in environments that are dedicated or mixed (with other services), you have the flexibility to choose your architecture to take advantage of Files services.

Nutanix Files streamlines design, implementation, and maintenance, providing an unrivaled experience for key use cases like end-user computing, video capture, medical imaging archives, application data, and the many other projects that store data in shared environments.

Appendix

About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix enterprise cloud software leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization, and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications. Learn more at www.nutanix.com or follow us on Twitter [@nutanix](https://twitter.com/nutanix).

List of Figures

Figure 1: Nutanix Enterprise Cloud OS Stack.....	9
Figure 2: Nutanix Files Instances Run as VMs.....	11
Figure 3: Data Path Architecture of Nutanix Files.....	12
Figure 4: FSVM Internal Communication on One Node.....	13
Figure 5: FSVM vDisks and Volume Groups.....	14
Figure 6: Initial 40 TB File System Pool.....	15
Figure 7: 80 TB File System Pool Following First Expansion.....	16
Figure 8: 120 TB File System Pool After Second Expansion.....	16
Figure 9: 140 TB File System Pool After Final Expansion.....	16
Figure 10: Distributed Directory Shares.....	17
Figure 11: Two Standard Shares on the Same File Server.....	19
Figure 12: Create Share Menu with Nested Share Path.....	22
Figure 13: Submounting a Distributed Share into a Standard Share.....	23
Figure 14: Distributed Share and Top-Level Directories.....	24
Figure 15: Three Standard Shares.....	24
Figure 16: Performance Optimization Recommendation.....	26
Figure 17: Load Balancing Volume Groups with Nutanix Volumes.....	27
Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups.....	28
Figure 19: Each FSVM Controls Its Own Volume Groups.....	29
Figure 20: FSVM-1 Failure.....	30
Figure 21: DNS Request for SMB.....	33
Figure 22: Shared Folders MMC.....	35
Figure 23: Enable SMB Encryption.....	37

Figure 24: Nutanix Files User Management with NFS.....	39
Figure 25: DNS Request for NFSv4.....	40
Figure 26: DNS Request for NFSv3.....	41
Figure 27: Multiprotocol Shares.....	42
Figure 28: Enable Non-Native Protocol Access.....	43
Figure 29: Multiprotocol User Mapping.....	44
Figure 30: Setting a Share Quota.....	46
Figure 31: Blocked File Types for the File Server.....	48
Figure 32: Blocked File Types for a Share.....	48
Figure 33: File Analytics Dashboard.....	50
Figure 34: Data Age Analytics.....	51
Figure 35: File Analytics Audit Trails.....	52
Figure 36: File Analytics Anomaly Detection.....	53
Figure 37: File Analytics Usage Anomalies Dashboard.....	54
Figure 38: File Analytics Ransomware Page.....	55
Figure 39: Report Builder.....	56
Figure 40: Saved Reports.....	56
Figure 41: Smart Tier Architecture.....	58
Figure 42: Data Lens Dashboard.....	59
Figure 43: Data Lens Age Explorer.....	60
Figure 44: Data Lens Tiering.....	61
Figure 45: Tiered File.....	62
Figure 46: Data Summary by Age and Storage Tier.....	63
Figure 47: Protection Using SSR.....	65
Figure 48: Nutanix Files Cloning Use Cases.....	68
Figure 49: File Manager in Prism Central.....	69

Figure 50: Files Manager Data Protection.....	69
Figure 51: Files Smart DR.....	70
Figure 52: Data Protection Policy.....	71
Figure 53: Data Protection Compliance.....	71
Figure 54: Files Smart DR Failover.....	72
Figure 55: ICAP Workflow.....	74
Figure 56: Configure Multiple ICAP Servers.....	75
Figure 57: Default Scan Settings.....	76
Figure 58: Quarantined Files.....	77
Figure 59: CFT Backup Process.....	80
Figure 60: Adding More Readers to Drive Throughput on a Home Share.....	82
Figure 61: Backup Job Details.....	83

List of Tables

Table 1: Document Version History.....	6
Table 2: Supported Active Client Connections.....	32
Table 3: User Mapping Requirements.....	44
Table 4: Order of Precedence for Quotas.....	47

How to buy

Contact



+1.876.931-9552

mybusiness@infoexchangeja.com