

Dell EMC Unity: NAS Capabilities

Abstract

This white paper explains the NAS capabilities available on Dell EMC Unity storage systems. It provides a detailed review of the functionality, features, and protocols supported, as well as enhancements enabled by the Dell EMC Unity file system architecture.

January 2021

Revisions

Date	Description
May 2016	Initial release – Unity OE 4.0
December 2016	Updated for Unity OE 4.1
July 2017	Updated for Unity OE 4.2
March 2018	Updated for Unity OE 4.3
August 2018	Updated for Unity OE 4.4
January 2019	Updated for Unity OE 4.5
June 2019	Updated for Unity OE 5.0
October 2019	Template and format update
July 2020	Minor updates
January 2021	Minor updates

Acknowledgements

Author: Wei.C.Chen@dell.com

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [2/3/2021] [Technical White Paper] [H15572.9]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	6
Audience	6
1 Introduction.....	7
1.1 Terminology	7
2 NAS servers	10
2.1 Interfaces	10
2.2 High availability.....	10
2.2.1 Link aggregation	10
2.2.2 Fail safe networking.....	11
2.3 Advanced static routing	11
2.4 Packet reflect	13
2.5 IP multi-tenancy	14
2.6 NAS server mobility	17
2.7 NAS parameters	18
3 Dell EMC Unity file system	19
3.1 Scalability.....	19
3.2 Storage efficiency	20
3.3 Availability and recoverability	20
3.4 Performance	20
3.5 Virtualization	20
3.6 File system attributes.....	23
4 Shrink and extend	26
4.1 Manual extension.....	26
4.2 Manual shrink	26
4.3 Automatic shrink	28
4.4 Automatic extension	29
4.5 Minimum allocation size	31
5 File-Level Retention	32
6 Quotas	33
6.1 Quota types	34
6.2 Quota limits.....	35

6.3	Quota policy	36
7	Protocol options.....	37
7.1	SMB	37
7.1.1	Sync writes enabled	38
7.1.2	Oplocks enabled.....	38
7.1.3	Notify on write/access enabled.....	39
7.1.4	Continuous availability.....	39
7.1.5	Protocol encryption	40
7.1.6	Access-based enumeration	40
7.1.7	BranchCache	40
7.1.8	Offline availability.....	41
7.2	NFS.....	41
7.2.1	Parameters	43
7.2.2	NFSv4.....	45
7.2.3	Secure NFS	45
7.2.4	vVols	46
7.2.5	Host access	46
7.3	Multiprotocol	47
7.3.1	Directory services	48
7.3.2	SMB	48
7.3.3	NFS.....	48
7.3.4	User mapping	53
7.3.5	Default accounts for unmapped users.....	54
7.3.6	Automatic mapping for unmapped Windows accounts	55
7.3.7	Mapping process	56
7.3.8	Mapping management and diagnostics.....	58
7.3.9	Additional options	59
7.3.10	Access policy	59
7.3.11	UMASK.....	61
7.4	Locking and folder rename policy.....	63
7.4.1	Locking policy	63
7.4.2	Folder rename policy	64
7.5	FTP and SFTP.....	65
7.6	Internationalization	67
8	Features	68
8.1	Data reduction	68

8.2	Local protection	69
8.2.1	Snapshots	69
8.2.2	NDMP	70
8.3	Remote protection	71
8.3.1	MetroSync	71
8.3.2	Asynchronous replication	72
8.3.3	RO proxy NAS servers	76
8.3.4	RW SMB proxy shares	78
8.3.5	Interfaces	80
8.4	FAST technology	80
8.5	Custom file alert thresholds	80
8.6	Common Event Enabler	81
8.6.1	CAVA	81
8.6.2	CEPA	82
8.7	Cloud Tiering Appliance	83
8.8	File import	85
9	Conclusion	87
A	Technical support and resources	88
A.1	Related resources	88

Executive summary

Dell EMC™ Unity storage sets the new standard for midrange storage with a powerful combination of simplicity, modern design, affordable price point, and deployment flexibility — perfect for resource-constrained IT professionals in large or small companies. It delivers a full block and file unified environment in a single 2U enclosure. The system allows using the same pool to provision and host LUNs, consistency groups, NAS servers, file systems, and VMware® vSphere® Virtual Volumes™ (vVols) alike. The Unisphere™ management interface offers a consistent look and feel whether managing block resources, file resources, or both.

The Dell EMC Unity file system is a 64-bit architecture introduced on the Dell EMC Unity family of storage systems. This file system architecture allows for unprecedented scalability, efficiency, and flexibility, as well as a rich set of features to allow file storage administrators to leverage Dell EMC Unity storage systems for a wide range of traditional and transactional NAS use cases. Whether configuring home directories or deploying performance-intensive applications on file storage, the Dell EMC Unity file system provides the feature set and deep virtualization integration necessary for any storage environment.

The Dell EMC Unity file system was designed to integrate seamlessly with Dell EMC Unity block storage through similar configuration and management workflows that greatly reduce the management overhead traditionally associated with file storage. Similarly, the architecture allows file and block to share the same pools and features, resulting in a truly unified offering in the storage market. Features such as data protection and storage efficiency behave uniformly across file and block storage resources and benefit both equally.

In addition, Dell EMC Unity storage offers advanced NAS capabilities to provide additional value. Dell EMC Unity systems are designed to operate in networking environments with multiple VLANs, subnets, and gateways by providing advanced static routing and packet reflect. They also support multiple tenants residing on the same system by separating network traffic at the kernel level to provide enhanced security and dedicated network resources to each tenant. Common event enabler allows applications to scan for viruses and receive file event notifications for auditing, quota management, search/indexing, and more. Cloud Tiering Appliance integration enables tiering to cloud repositories based on user-configured policies.

Unisphere provides a powerful unified management framework composed of an HTML5 graphical user interface, command line interface, and RESTful API allowing novice and experienced administrators alike to easily manage their file storage environments. Wizard-based file provisioning enables novice administrators to quickly get a file storage environment up and running. The CLI and RESTful API allow more seasoned administrators to create complex scripts to facilitate specific use cases, while still using the Unisphere GUI for daily provisioning and management tasks. Most importantly, file and block management functionality is available from within all interfaces, ensuring a uniform user experience regardless of the task. In addition, a Python StorOps storage management library and PowerShell cmdlets are available to manage Dell EMC Unity systems.

Audience

This white paper is intended for Dell EMC customers, partners, and employees who are interested in the file storage functionality on Dell EMC Unity storage systems. It is assumed that the reader is at least an IT generalist who has experience as a system or network administrator.

1 Introduction

Dell EMC Unity storage systems take a unique approach to file storage in that file is tightly integrated with block, resulting in a unified storage solution. Dell EMC Unity storage employs storage pools which are used for all resource types directly, meaning LUNs, file systems, and even VMware® vSphere® Virtual Volumes™ (vVols) can be provisioned out of the same unified pools. When provisioning file systems, administrators simply provision file systems as they would traditionally provision LUNs, by choosing a storage pool. Because Dell EMC Unity storage is truly unified in both its hardware and software architecture, there is no need for the additional management overhead of provisioning LUNs, presenting to an internal gateway, or creating file storage pools. This drastically simplifies management and also allows the system to leverage a core set of unified features for both block and file, since both types of storage are implemented and provisioned at the same level using the same hardware.




Figure 1 Unified storage pool

1.1 Terminology

Allocated space: The actual amount of capacity that is provisioned to a storage resource (such as a file system, LUN, or VMware datastore) from the storage pool, not including snapshots and thin clones. For thick provisioned storage resources, the allocated space is equal to the requested capacity. For thin provisioned storage resources, the allocated space is the capacity that is currently provisioned from the storage pool, which could be less than the requested capacity of the storage resource.

File system: A storage resource that can be accessed through file sharing protocols such as SMB or NFS.

Fully Automated Storage Tiering for Virtual Pools (FAST™ VP): A feature that relocates data to the most appropriate disk type depending on activity level to improve performance while reducing cost.

FAST Cache: A feature that allows Flash disks to be configured as a large capacity secondary cache for the Pools on the system.

NAS server: A Dell EMC Unity storage server that uses the SMB, NFS, or FTP/SFTP protocols to catalog, organize, and transfer files within designated file system shares. A NAS server, the basis for multi-tenancy, must be created before you can create file-level storage resources such as file systems or VMware file datastores.

Network File System (NFS): An access protocol that enables users to access files and folders located on a network. Typically used by Linux®/Unix hosts.

Non-base allocated space: The amount of pool space used for the snapshot and thin clones, if applicable. This is displayed at the pool level.

Oversubscription: A storage provisioning method that allows administrators to provision more capacity than may be physically available in a particular storage pool. When thin provisioned storage resources are associated with a common storage pool, they can potentially request (or subscribe to) more storage capacity than the storage pool contains. Administrators can then add more drives to the system or assign more drives to the storage pool as needed. Hosts connected to thin provisioned storage resources are unaware of the pool oversubscription. They see the subscribed (or maximum) size for each thin provisioned storage resource, not the current allocated size.

Server Message Block (SMB): An access protocol that allows remote file data access from clients to hosts located on a network. This is typically used in Microsoft® Windows® environments.

Size: The client visible size of a storage resource, as set at the time of creation or afterward, regardless of the actual amount of space consumed by the storage resource from the pool (see Total Pool Space Used). Size may be larger than the actual allocated size for thinly provisioned storage resources, forming the basis for overprovisioning.

Snapshot: A point-in-time view of data stored on a storage resource. A user can recover files from a snapshot, restore a storage resource from a snapshot, or provide access to a host. Snapshots can be read-only or read/write.

Storage pool: A collection of disk drives configured with a particular storage profile. The storage profile defines the type of disks used to provide storage and the type of RAID configured on the disks. The storage pool's configuration defines the number of disks and quantity of storage associated with the pool. Dell EMC Unity uses unified storage pools for both block and file storage resources.

Storage processor (SP): A storage node that provides the processing resources for performing storage operations as well as servicing I/O between storage and hosts.

Thin provisioned storage resource: A storage resource (such as a file system, LUN, or VMware datastore) that is not fully allocated from the storage pool. The client can see the full size of the storage resource even though only a portion of the storage resource is allocated from the storage pool.

Total pool space used: The total amount of space consumed by the storage resource on the pool, including all overhead, metadata, snapshots, and thin clones. This is displayed at the pool level.

Unisphere CLI (UEMCLI): The command-line interface for managing Dell EMC Unity storage systems.

Unisphere: The HTML5 web-based user interface for managing Dell EMC Unity storage systems.

Used space: The amount of space in a file system that is actually consumed by the clients. This relates to the amount of data users have stored in the file system.

Virtual Volumes (vVols): A VMware storage framework which allows VM data to be stored on individual volumes. This allows for features such as snapshots to be applied at a VM-granularity and provides Storage Policy Based Management (SPBM).

VMware vSphere Storage APIs Array Integration (VAAI): A set of APIs to enable communication between VMware vSphere ESXi™ hosts and storage devices. The APIs define a set of storage primitives that enable the ESXi host to offload certain storage operations to the array, which reduces resource overhead on the ESXi hosts and can significantly improve performance for storage-intensive operations such as storage cloning, zeroing, and so on. The goal of VAAI is to help storage vendors provide hardware assistance to speed up VMware I/O operations that are more efficiently accomplished in the storage hardware.

2 NAS servers

Because Dell EMC Unity has a single-enclosure, two-storage-processor architecture with no concept of designated file hardware, file data is served through virtual file servers known as NAS servers, which may reside on either storage processor. A NAS server, which is required before creating file systems, allows for basic multi-tenancy in that each contains its own distinct set of configuration information and file interfaces. Because each NAS server is logically separate, clients of one NAS server cannot access data on another NAS server and conversely. Each NAS server may contain up to 50 production and 10 backup interfaces and a variety of configuration information including naming services, sharing protocols, Microsoft Active Directory® domain settings, UNIX directory service, user mapping configuration, data protection settings and more. Once a NAS server with the appropriate protocol configuration exists, administrators can create file systems and leverage many of their advanced capabilities available on Dell EMC Unity storage.

2.1 Interfaces

Starting with Dell EMC Unity OE version 4.4, ports can be configured for a custom MTU size between 1280 and 9216. Previously, the MTU sizes were limited to either 1500 or 9000. The custom MTU size can be configured on ports that are used for NAS server, replication, and import interfaces. Any ports that have iSCSI interfaces created must still use 1500 or 9000. Ports with custom MTU sizes configured can be used for link aggregations and fail safe networking (FSN) as long as the MTU size matches on all ports. This feature enables Dell EMC Unity systems to be used in complex environments where customized MTU sizes are required.

2.2 High availability

On Dell EMC Unity systems, both SPs can be used simultaneously so no dedicated standby hardware is required. The peer SP acts as a hot standby, which actively services I/O but is also ready to take over additional resources if necessary. For example, if SPA fails, the NAS servers along with their file systems fail over to SPB. There may be a short interruption to host access during this operation.

In file environments, NAS servers include one or more network interfaces that are created on one or more Ethernet ports for host access. Link loss can be caused by many environmental factors such as cable or switch port failure. In case of link loss, the system does not initiate a failover of the NAS server to the peer SP. Therefore, it is important to configure high availability on the ports to protect against these types of failure scenarios.

2.2.1 Link aggregation

Link aggregation combines multiple physical network connections into one logical link. This provides increased bandwidth by distributing traffic across multiple connections and also provides redundancy in case one or multiple connections fail, depending on the configuration. If connection loss is detected, the link is immediately disabled, and traffic is automatically moved to the surviving links in the aggregate to avoid disruption. The switch should be properly configured to add the ports back to the aggregate when the connection is restored. Although link aggregations provide more overall bandwidth, each individual client still runs through a single port. Dell EMC Unity systems use the Link Aggregation Control Protocol (LACP) IEEE 802.3ad standard.

Link aggregations can be configured with two to four ports. Starting with Dell EMC Unity OE version 4.2.1, link aggregation can be created using ports from different I/O modules and also between I/O modules and the on-board Ethernet ports. Previously, only ports belonging to the same I/O module or on-board Ethernet ports

could be aggregated together. All ports within the aggregation must have the same speed, duplex settings, and MTU size.

Link aggregation can be used for NAS server, replication, and file import interfaces. Link aggregation is not supported for iSCSI since multipathing is used for block access. Any ports that have iSCSI interfaces created on them are not listed as options when creating a link aggregation. Also, link aggregations devices are not listed as options when creating iSCSI interfaces.

2.2.2 Fail safe networking

Starting with, Dell EMC Unity OE version 4.2.1, fail safe networking (FSN) is available. FSN is a high availability feature that extends link failover into the network by supporting switch-level redundancy. FSN appears as a single link with a single MAC address and potentially multiple IP addresses. FSN can consist of Ethernet ports, link aggregations, or any combination of the two. FSN can be created using ports from different I/O modules and also between I/O modules and the on-board Ethernet ports.

FSN adds an extra layer of availability to link aggregations alone as link aggregations provide availability in the event of a port failure while FSN provides availability in the event of a switch failure. Each port or link aggregation is considered as a single connection and only the primary port or link aggregation in an FSN is active at a time. All ports in an FSN must have the same MTU size, but the speed and duplex settings can vary.

If the system detects a failure of the active connection, it automatically switches to the standby connection in the FSN. That new connection assumes the network identity of the failed connection, until the primary connection is available again. You can designate which connection is the primary connection at creation time. To ensure connectivity in the event of a hardware failure, create FSN devices on multiple I/O modules or on-board ports. The FSN components can be connected to different switches and no special switches are required. If the network switch for the active connection fails, the FSN fails over to a connection using a different switch, thus extending link failover out into the network.

For more information about high availability and redundancy, reference the Dell EMC Unity: High Availability white paper on [Dell EMC Support](#).

2.3 Advanced static routing

On a NAS server, interfaces can be configured to enable communication between the NAS server, client, and external services. The system automatically creates a local route when an interface is created. This directs traffic to the local subnet through the interface that's local to that subnet. If a default gateway is entered, a default route is also created. This directs all non-local traffic to the default gateway, which forwards it to other networks. In addition to these system-created routes, user-defined static routes can also be created.

In addition, starting with Dell EMC Unity OE version 4.1, static routes can also be configured to determine where to forward a packet so that it can reach its destination. Static routes can be configured for both IPv4 and IPv6 interfaces. Each NAS server interface has its own independent routing table with up to 20 routes.


Using static routes enables the NAS server to access a destination using a specific gateway and interface. For example, complex networking environments may leverage multiple gateways, with each gateway enabling access to a different subnet. In this scenario, static routes must be configured to ensure packets are sent to the correct gateway for each subnet, instead of using a default gateway.

Static routes can either be a host or network route. A host route is the most specific type of route, which is only used when traffic is sent to a specific IP address. A network route is less specific, and is used when

sending traffic to a specific subnet. The system uses the most specific route available. If no host or network routes are defined, the default route is used (if configured).

New routes can be configured in the following pages in Unisphere:

- **NAS server Properties → Network → Interfaces and Routes:** The per-interface routing table can be displayed and managed by selecting an interface and clicking Show external routes for interfaces. This displays the list of routes that are used for inbound connections, such as communication initiated by a client for I/O.
- **NAS server Properties → Network → Routes to External Services:** The NAS server routing table is used for outbound communication initiated by the NAS server to external services, such as LDAP or DNS. This table is dynamically created by merging the per-interface routing tables to support dynamic interface configuration and replication. It ensures that the best possible routing configuration is used by the NAS server when interfaces get added, deleted or edited, either manually or due to the replication status changes.
- **Settings → Access → Routing:** Allows for the viewing and management of all static routes configured on the entire system.



NASServer Properties

General | **Network** | Naming Services | Sharing Protocols | Protection & Events | Security | Replication

Interfaces & Routes

• **Routes to External Services**

External Services Access Routes: ⓘ The following routes are used for outbound connections.

		From	Gateway	Destination	Subnet Mask / Prefi...	Replication Sync
<input type="checkbox"/>	!	10.0.0.1	10.0.0.254	10.0.1.1	255.255.255.255	
<input type="checkbox"/>	✓	10.0.0.2	Local	10.0.0.0	255.255.255.0	
<input type="checkbox"/>	✓	10.0.0.2	10.0.0.253	10.0.2.0	255.255.255.0	
<input type="checkbox"/>	✓	10.0.0.2	10.0.0.254	Default	0.0.0.0	
<input type="checkbox"/>	✓	10.0.5.1	Local	10.0.5.0	255.255.255.0	

Close

Figure 2 Routes to External Services

Enter the following information to create a new route:

- **From:** Select the interface for the route
- **Type:** default, host or net
- **Gateway:** Router on the local subnet to send this traffic to
- **Destination:** Destination IP address or network address
- **Subnet Mask / Prefix Length:** Destination network subnet mask or prefix length (network routes only)

Figure 3 New route

2.4 Packet reflect

Packet reflect, available starting in Dell EMC Unity OE version 4.1, is a feature that ensures outbound (reply) packets are sent back to the same host or router as the inbound (request) packet. This enables the NAS server to bypass routing and ARP table lookups when replying to a packet, so no routing configuration is required. With packet reflect, information including the local IP, remote IP, and next-hop MAC address are cached from the incoming packet. When the NAS server replies to that packet, it leverages this information to send the outbound packet to the proper location.

With packet reflect enabled, reply packets are always returned to the local MAC interface from which the request packet was sent, regardless of the destination IP address. For example, if a packet is received from a local gateway that is unknown to the NAS server's routing table, the reply packet will be returned to that unknown gateway, independent of the destination IP address. The return path is not influenced by the routing table. However, with packet reflect disabled, the reply packet path is determined by the destination IP address and the routing table. In this scenario, reply packets are returned using the routes defined in the routing table. The chosen path could be different than that of the originating unknown local gateway. This feature can be disabled (default) or enabled at the NAS server level and takes effect immediately.

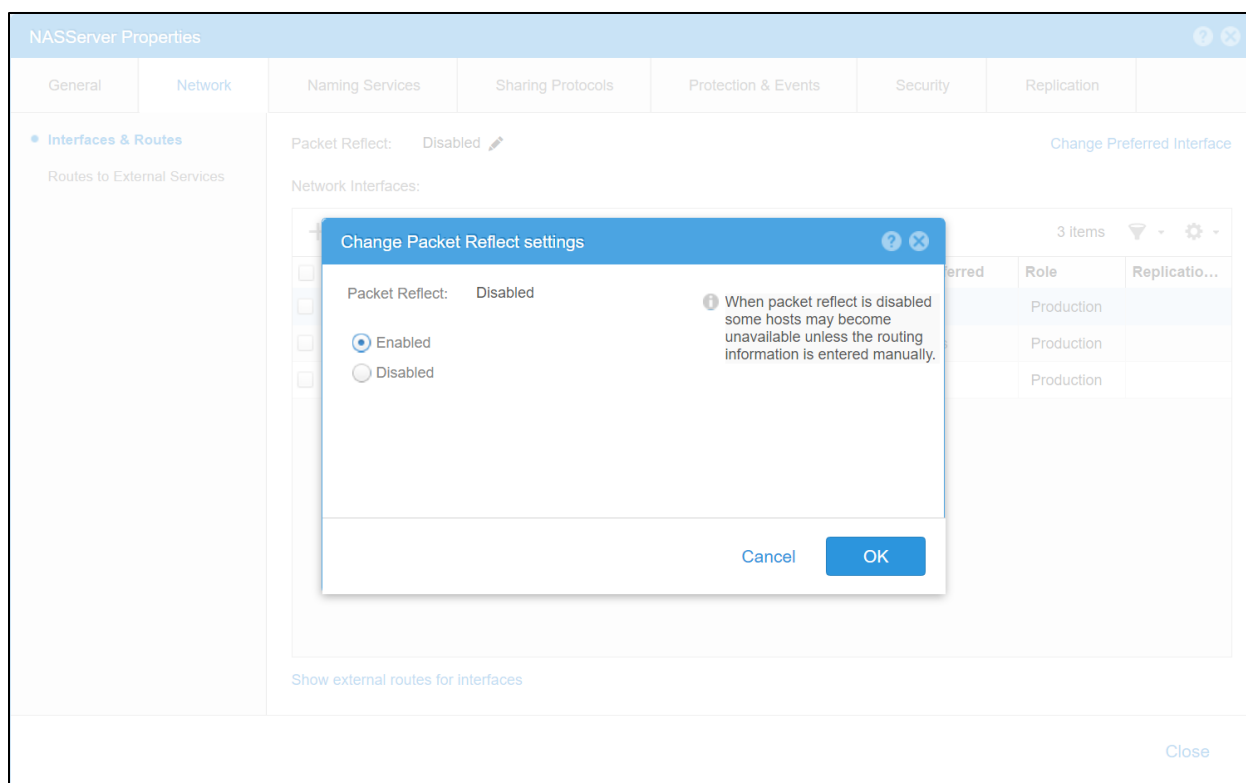


Figure 4 Packet reflect

Packet reflect enables dynamic routing automatic configuration and does not require any changes to the infrastructure. An example is if there is a router failure, replacement, or IP change. In these scenarios, packets may still be able to reach the Dell EMC Unity system through a different router. Packet reflect enables the Dell EMC Unity system to seamlessly adapt by sending the reply packets to the new router, since that is where the request packet was received from. This is an advantage compared to static routes, which must be manually managed by the network administrator.

Although Packet Reflect works for the majority of the communication to a NAS server, such as client-initiated file system IO, it is important to note that it does not work for communication initiated by the NAS server. For example, communication to a DNS or LDAP server still requires routing and ARP table lookups since there is no inbound packet to cache the necessary information from. Regardless of whether or not Packet Reflect is enabled, it is important to configure the appropriate routes on the **Routes to External Services** page to allow the NAS server to access these services.

2.5 IP multi-tenancy

Dell EMC Unity supports the ability to host multiple tenants on a single system, such as for service providers. Each NAS server has its own independent configuration which can be tailored to each tenant's requirements. File systems cannot be accessed from any other NAS server, other than the one that they are associated with. Dell EMC Unity's flexible pool-based architecture also enables the ability to separate each tenant onto their own pools for separation at the physical drive level, if necessary.

Dell EMC Unity OE version 4.1 includes support for IP multi-tenancy, which adds the ability to provide network isolation for tenants. This feature separates network traffic at the kernel level on the SP, enabling the ability to provide dedicated network resources for each tenant. Each tenant has its own dedicated network namespace including VLAN domain, routing table, firewall, interfaces, DNS and more. This also enables the

ability for multiple tenants to use the same or overlapping IP network configuration, so IPs can be duplicated across tenants. This avoids network interference between tenants and also enhances security. The separate network VLAN can only be maintained if the routers connected to those VLANs are also separate, or if there is a common router, it must have separated or partitioned routing tables that do not route across tenant's VLANs. This feature is only available on purpose-built Dell EMC Unity systems and is not available on Dell EMC UnityVSA.

In order to leverage IP multi-tenancy, switches need to be configured for VLAN tagging. Once VLANs are configured, tenant objects must be created in the **File → Tenants** page. When creating a tenant, enter the following information:

- **Name:** Tenant Name
- **UUID (Universal Unique Tenant ID) (optional):** Use the system generated UUID or enter it manually for existing tenants (for example, on a replication target system)
- **VLANs:** Select the VLAN(s) that are associated with this tenant

The screenshot shows the 'Add Tenant' dialog box. The 'Name' field is filled with 'Tenant_Finance'. The 'Enter UUID manually' checkbox is unchecked. The 'VLANs' list box contains the value '500'. The 'Add' and 'Remove' buttons are visible next to the list box. The 'Cancel' and 'OK' buttons are at the bottom right.

Figure 5 Create tenant

After tenants are created, create NAS servers and hosts (for NFS access) and associate them with the appropriate tenant. By default, NAS servers and hosts do not have any tenant association so one must be assigned if you want to use this feature. Each NAS server or host can only be associated with a single tenant and this can only be done at creation. After creation, the tenant configuration cannot be modified in any way. This is intentionally prohibited for security purposes.

If a NAS server is created with a tenant association, its interfaces must be created on one of the VLANs that's assigned to the tenant. For example, if Tenant_Finance has VLANs 500 and 501 assigned, interfaces on any NAS servers associated with this tenant must reside on these VLANs. Each VLAN can only be associated with one tenant at a time, but the assigned VLANs for each tenant can be modified at any time. The system ensures each tenant has unique VLAN assignments to provide isolation from other tenants. It is important to

note that the network infrastructure must also be configured with the appropriate VLANs to enable communication.

The screenshot shows a web-based configuration interface for creating a NAS server. The title bar reads 'Create a NAS Server'. On the left is a sidebar with navigation links: General (selected), Interface, Sharing Protocols, Unix Directory Service, DNS, Replication, Summary, and Results. The main area is titled 'Configure NAS Server General Settings'. It contains four fields: 'Server Name' with the value 'NASServer', 'Tenant' with a dropdown menu showing 'Select or enter a tenant.' and 'Tenant_Finance' selected, 'Pool' with the value 'Tenant_Finance', and 'Storage Processor' with an empty dropdown. At the bottom right are 'Cancel' and 'Next' buttons.

Figure 6 NAS server tenant association

After IP multi-tenancy is configured, the system can be configured to use duplicate IPs across multiple tenants. This enables each tenant to use any IP schema they wish, without worrying about interference with other tenants. Although multiple tenants may share the same IP address, they remain separated since they are on different VLANs. Because of this, each tenant can only access the NAS servers that are assigned to their tenant.

If a single tenant has multiple NAS servers, all NAS server interfaces must still be unique since it is within the same IP namespace. This feature only enables duplicating IP addresses with other tenants. If IP multi-tenancy is not used, all interfaces on the entire system must be unique since the default IP namespace is shared across the entire system.

With IP multi-tenancy enabled, external services such as DNS, LDAP, or NIS can be dedicated or shared. If each tenant provides their own external services, each NAS server can be configured to use these dedicated servers. However, for tenants that do not provide this, their NAS servers can also be configured to use the service provider's shared external services. This provides additional flexibility by allowing each tenant to configure external services depending on their use case.

If a NAS server is being replicated, the destination NAS server must have a matching tenant configuration. For example, you cannot replicate a non-tenanted NAS server to a tenanted NAS server. Tenants must be created on the target system utilizing the same UUID as on the source system.

Also, note that IP multi-tenancy only supports file systems. Using NFS datastores and vVols on a NAS server with IP multi-tenancy enabled is not supported. You are prohibited from enabling the vVol protocol endpoint

on a NAS server that has a tenant association. In order to use vVols, you must use a NAS server that does not have a tenant assigned.

On Dell EMC Unity OE version 4.1, total bandwidth historical metrics are available at a tenant-level granularity. This provides the total amount of I/O requests in KB/s for the selected tenant. Dell EMC Unity OE version 4.2 also adds real-time read/write bandwidth metrics at a tenant-level granularity. This displays the amount of read or write I/O requests, in KB/s, for the selected tenant.

2.6 NAS server mobility

When creating a new NAS server, there is an option to choose the SP owner which applies to the NAS server and its associated file systems. Dell EMC Unity OE version 4.2 introduces the ability to change the SP owner on existing NAS server after they are created. This is automatically applied to all of the NAS server's associated file systems, enabling this feature to be used for load balancing purposes or to resolve network issues. The figure below shows how to change the SP owner of a NAS server.

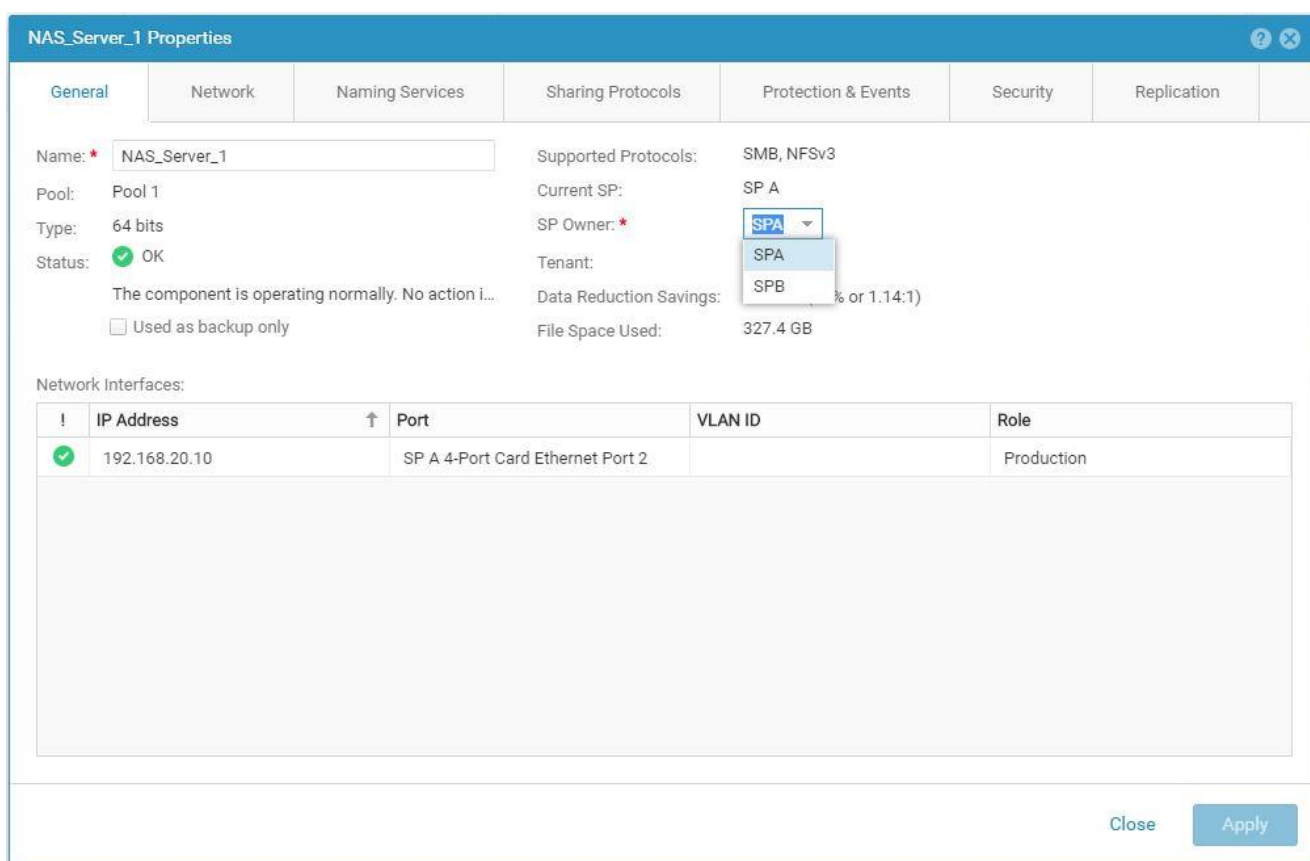


Figure 7 NAS server mobility


If a NAS server is moved to the peer SP, all of its existing configuration and features are carried over and continue to work as normal, such as CAVA, CEPA, VLANs, DNS, NIS, LDAP, and so on. Replication sessions also have the ability to move with the NAS server. In order to do this, all replication sessions must be in a paused state prior to starting the mobility operation. If a NAS server mobility operation is attempted while replication sessions are still active, an error is returned. While the mobility operation is in progress, all replication commands, besides show to view the session details, are rejected. Once the mobility operation is complete, the replication sessions can be resumed. Group pause and resume operations are also available in

Dell EMC Unity OE version 4.2 which simplifies this process. For more information on group operations, refer to section 8.3 on remote protection.

2.7 NAS parameters

Parameters are used for controlling the desired behavior and advanced tweaking of NAS related features. All parameter changes are preserved through SP reboots and NAS server mobility operations. Management of parameters is only available in the CLI and can be managed by the service user. As of Dell EMC Unity OE version 4.1, the ability to configure system parameters is available. Dell EMC Unity OE version 4.2 includes an additional enhancement which enables parameters to be configured at a NAS server level. This enables a higher level of granularity and also allows parameters to be automatically replicated as part of NAS server replication. As part of this enhancement, the `svc_param` command is deprecated and is replaced by the `svc_nas -param` command on Dell EMC Unity OE version 4.2 and later.

When viewing details or changing a parameter, the output provides additional information about the parameter including the granularity and when it takes effect. Certain parameters are still global, which means they can only be applied at the system level. Also, some parameter changes require a SP or NAS server reboot to take effect. The figure below shows an example of a NAS server parameter.



```

18:54:08 service@none spa:~> svc_nas NASServer -param -f cvfs -i virtualDirName

NASServer :
name                = virtualDirName
facility_name        = cvfs
default_value        = ckpt
current_value        = ckpt
configured_value     =
param_type           = global
user_action          = reboot SP
change_effective     = reboot SP
range                = *
description          = Defines virtual directory name used to traverse mounte
d checkpoints

18:54:19 service@none spa:~>

```

Figure 8 NAS server parameters

As part of this feature, the ability to restart an individual NAS server is introduced so parameter changes can be applied without affecting other NAS servers on the system. Restarting a NAS server causes a brief outage to clients that are not running SMB3+CA or NFS.

For more information about all of the available NAS server parameters and how to configure them, reference the *Service Commands* document on [Dell EMC Support](#).

3 Dell EMC Unity file system

The Dell EMC Unity file system brings a number of improvements over existing NAS file system technologies. With the 64-bit architecture, the Dell EMC Unity file system is able to scale far beyond the limitations of previous file systems in many areas, including file system size. The file system is also flexible and well suited to both traditional and transactional use cases, providing value over existing technologies in a variety of ways including:

- Scalability
- Storage Efficiency
- Availability and Recoverability
- Performance
- Virtualization

In addition, Dell EMC Unity file systems include a full set of features, enabling them to be utilized and protected as efficiently as possible. While several features including quotas, shrink, and reclaim are purpose built for Dell EMC Unity file systems, others leverage Dell EMC Unity's deep integration between block and file to provide truly unified features applicable to both block and file storage resources.

3.1 Scalability

Dell EMC Unity file systems allow for enhanced scalability in a number of different areas, including maximum file system size. Dell EMC Unity file systems can accommodate more data, directories, and files than previous file system architectures, making Dell EMC Unity ideal for traditional and transactional NAS use cases. The table below covers several of the scalability attributes of file systems in Dell EMC Unity.

Table 1 File system scalability

File system attribute	Dell EMC Unity file system
Maximum File System Size	256 TB (Dell EMC Unity OE version 4.2+)
Subdirectories per Directory	~10 million
Files per File System	~32 billion
Filenames per Directory	~10 million
ACL IDs	4 million
Timestamp Granularity	1 nanosecond

As of Dell EMC Unity OE version 4.2, the maximum file system size is increased from 64 TB to 256 TB for all file systems, including existing file systems created on earlier code. File system sizes can range from 3 GB to 256 TB. Any attempts to create a file system with a size outside of this range results in an error.

Administrators also have the ability to shrink and extend file systems to any size within the supported limits. Clients have the ability to fill up the file system with up to 256 TB of data. It is important to note that larger file system sizes result in longer backup times.

All thin file systems, regardless of size, have 1.5 GB reserved for metadata upon creation. For example, after creating a 100 GB thin file system, Unisphere immediately shows 1.5 GB used. When the file system is

mounted to a host, it shows 98.5 GB of usable capacity. This is because the metadata space is reserved from the usable file system capacity. This avoids situations where creating a 100 GB file system actually consumes 101.5 GB of capacity out of the storage pool.

3.2 Storage efficiency

Dell EMC Unity's unique file system architecture and unified storage pools allow for extreme flexibility as changes arise in file storage environments. File systems provisioned in Unisphere are thin by default, which is required for data reduction. Also, starting with Dell EMC Unity OE version 4.2, Unisphere also provides the ability to create thick file systems. Previously, thick file systems could only be created using CLI. There is no way to convert from thin to thick or thick to thin after a file system is created.

File systems can easily be extended to provide more capacity or shrunk to reclaim unused space back to the unified pool to be available for use by any type of resource. Dell EMC Unity also intelligently monitors thin file systems continuously for suboptimal space utilization, and initiates automatic extension and shrink operations as needed to ensure capacity is being used as efficiently as possible. These operations are fully integrated with all of Dell EMC Unity's features to ensure that file system size can always be modified to best fit changing environments without impacting or being restricted by data protection or performance requirements.

3.3 Availability and recoverability

Dell EMC Unity file systems include enhanced availability and recoverability measures in order to minimize downtime. Fault containment and panic avoidance allows the Dell EMC Unity system to recover corrupted file systems while they remain online in some cases, and avoids impacting the file system's associated NAS server in the case where a corrupted file system must be taken offline for recovery. Due to Dell EMC Unity's truly unified architecture, a file system does not share a second level "file pool" with other file systems. This means that there is no ability for a faulted file pool LUN to potentially affect multiple associated file systems, improving fault isolation.

3.4 Performance

The Dell EMC Unity file system is an entirely new file system architecture designed with both transactional and traditional NAS use cases in mind. Because of this, performance is a main priority, even in the presence of extreme scalability. Dell EMC Unity file systems are able to scale to maximum size without significant performance degradation, all while leveraging the multicore optimized architecture of Dell EMC Unity storage systems. For more information on best practices when configuring Dell EMC Unity file systems, reference the *Dell EMC Unity: Best Practices Guide* on [Dell EMC Support](#).

3.5 Virtualization

Dell EMC Unity also includes tight integration with VMware vSphere that benefits file storage administrators and virtualization administrators alike. In addition to traditional SMB and NFS file systems, Dell EMC Unity allows users to create a special NFS file system type optimized for VMware use. In Unisphere, this can be accomplished by creating a VMware NFS datastore from the VMware Storage page. When giving access to an ESXi host previously discovered from the VMware Access page, the VMware NFS datastore will be automatically detected and mounted as a datastore on the ESXi host with no manual intervention necessary.

In addition, Dell EMC Unity NFS VMware datastores give administrators the unique ability to select the underlying file system block size to best match the host I/O size of the intended application. A file system

block size is the smallest guaranteed physical mapping within a file system, which is set at 8 KB for Dell EMC Unity SMB and NFS file systems. However, because NFS datastores are often intended for specific application workloads, Dell EMC Unity provides the ability to set this block size to 8 KB, 16 KB, 32 KB, or 64 KB during datastore configuration to best accommodate the I/O size typically used by particular applications. Because administrators may not always be aware of the host I/O size of their intended application, Dell EMC Unity maintains an internal mapping of application to I/O size for popular applications, which allows users to simply specify the intended application instead. Dell EMC Unity will then configure the backend file system block size to match the I/O size used by this application. Applications with predefined host I/O sizes include:

- Exchange 2007
- Exchange 2010
- Exchange 2013
- Oracle
- SQL Server
- VMware Horizon VDI
- SharePoint
- SAP

The screenshot shows the 'Create VMware Datastore' wizard with the 'Storage' step selected. The 'Host IO Size' dropdown menu is open, showing a list of predefined application mappings and block sizes. The '8K (default)' option is highlighted.

Host IO Size
8K (default)
16K
32K
64K
Exchange 2007
Exchange 2010
Exchange 2013
Oracle
SQL Server
VMware Horizon VDI
SharePoint
SAP

Figure 9 Host IO Size

This approach of allowing administrators to specify the file system block size has two advantages over a single fixed block size. An 8 KB block size is unnecessarily granular for applications that address storage in

larger increments such as 64 KB, so it is more performance-efficient to match the file system block size to the application I/O size. Also, from a recovery perspective, fewer larger blocks reduce file system check (FSCK) times considerably when compared to more numerous smaller blocks. This is especially important for scaling purposes to avoid long FSCK times in the presence of the very large file systems supported on Dell EMC Unity. While this feature provides potential benefits, it is important to be sure of the correct application or I/O size setting when changing the default of 8 KB. Choosing an incorrect IO size can be detrimental to the performance of the file system, impose unnecessary flash wear penalties, and increase space consumption on small files when the configured IO size is larger than the actual IO size of the application. Because of this, it is recommended to leave the default and minimum IO size of 8 KB for general purpose VMware datastores or for those where the intended application of host IO size is unknown.

As an additional point of integration, Dell EMC Unity file systems support VMware vSphere Storage APIs Array Integration (VAAI) through a VAAI plugin, allowing for hardware acceleration and offloading through access to supported file primitives including FAST clone, snap-of-snap, extended statistics, and reserve space. Because of Dell EMC Unity's scalable 64-bit file system architecture, up to 256 levels of VM snapshots are possible with VAAI. With this capability, administrators can provision multiple levels of snapshots (also called fast clones) from a single golden image. In the figure below, a base VM is used as a source or golden copy for a snapshot to be taken. Similarly, a snapshot can then be taken of this snapshot. This process can then continue to create additional levels of hierarchical snapshots as necessary.

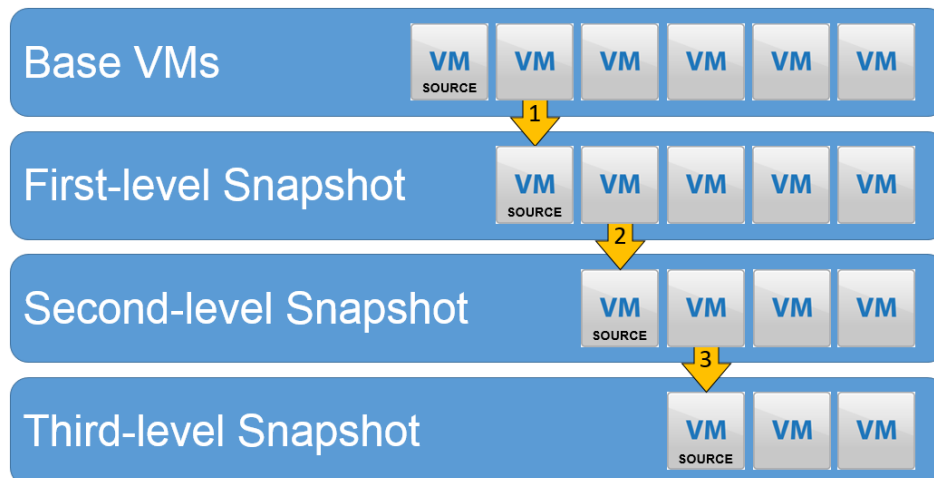


Figure 10 VAAI snap-of-snap

This functionality can be useful in many different cases, such as virtual desktop infrastructure or test and development. For example, these types of hierarchical snapshots can be leveraged as part of a software development environment where developers need to test out incremental changes to a base version of an operating system. As updates of minor software patches are installed, the test environment virtual machine could be snapped at every level in order to test the impact of each level of incremental software changes. This incremental testing and protection can continue until a final production version of the software is developed.

3.6 File system attributes

All file systems have the following properties related to capacity:

- **Size:** The provisioned capacity of the file system that is exposed to the client. For example, a newly created 256 TB thin file system has a client visible size of 256 TB, even though very little space is actually being consumed from the pool at this point.
- **Used space:** The amount of capacity consumed on the file system by clients. Prior to Dell EMC Unity OE version 4.5, the number reported by Unisphere and the number seen on the client may be slightly different. In Dell EMC Unity OE version 4.5 and later, the values in Unisphere will match what's seen on the client.
- **Allocated space:** The amount of space consumed from the pool for just the file system.
- **Total pool space used:** The total capacity consumed from the pool for the file system including the allocated space, snapshot allocation, metadata, and other overhead. Prior to Dell EMC Unity OE version 4.5, this attribute was only displayed in the Pool Properties page. Starting in Dell EMC Unity OE version 4.5, this attribute is also available in the File System Properties page.
- **Preallocated:** The amount of space that has been reserved for the file system but has not been used or reclaimed. Examples of preallocated include space reserved for new incoming data or space to be reclaimed after deleting a snapshot. Prior to Dell EMC Unity OE version 4.5, the preallocated space was included in the Total Pool Space Used number. Starting with Dell EMC Unity OE version 4.5, the preallocated space is displayed separately and is no longer part of the Total Pool Space Used number.

Prior to Dell EMC Unity OE version 4.5, the size and used space reported by Unisphere did not always match exactly what was displayed on the client. Starting with Dell EMC Unity OE version 4.5, an enhancement was added to ensure that Size and Used Space are consistent between Unisphere and the client.

Total pool space used should be used as the reference point when tracking pool capacity utilization and planning for future expansion. Size and total pool space used are related to overprovisioning. The sum of the sizes of all file systems in a pool may exceed the actual size of the pool, as long as the sum of the total pool space used does not exceed the space available in the pool. For example, a 10 TB pool may contain six 2 TB file systems, as long as the sum of the file systems' total pool space used does not exceed 10 TB.

Used space is the actual capacity consumed by the clients on the file system. For example, a 2 TB file system has 500 GB of files actually residing on it. The file system may have 600 GB allocated to it and consume a total of 625 GB of space from the storage pool after accounting for preallocation, snapshots, metadata, and overhead. In this case, size = 2 TB, Allocated space = 600 GB, Total pool space used = 625 GB, and used space = 500 GB. This is shown in the figure below.

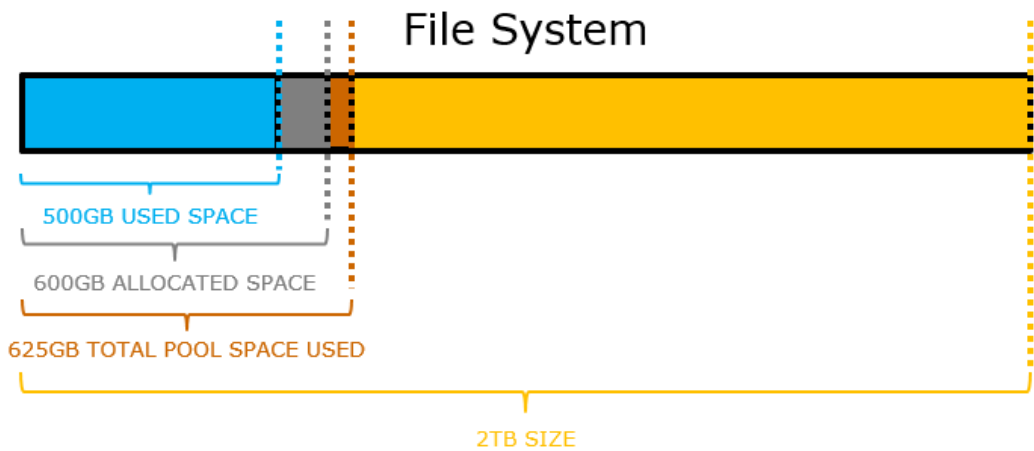


Figure 11 File system space

Generally, size ≥ allocated space ≥ used space. However, if the file system contains sparse files or if data reduction is enabled, it is possible for used space to be greater than allocated space. Data reduction reduces the amount of physical storage needed to store the client data on the storage pool. For example, a file system can have size = 2 TB, allocated space = 100 GB, total pool space used = 125 GB, used space = 500 GB, and data reduction savings = 350 GB. In this case, the file system reports 25% full since 500 GB / 2 TB is used on the file system, but only 125 GB is actually consumed from the pool due to the data reduction savings. The client does not see any difference from the host point of view on the file system, but the saved storage can be used for other resources on the system.

The numbers shown in these examples are for illustration purposes only. The actual amount of capacity used and saved will vary depending on the configuration and workload. For more information about on data reduction works, reference the *Dell EMC Unity: Data Reduction* white paper on [Dell EMC Support](#).

As capacity is consumed on the file system, additional capacity is allocated from the pool. This continuously happens until the size is reached and the file system becomes full. Capacity information is shown in Unisphere on the file system properties page to illustrate file system space utilization. The figure below shows how file systems might appear in Unisphere, where the size is shown as a numeric value while the allocated and used space are shown using bars similar to those on the previous figure.

File Systems								
SMB Shares NFS Shares NAS Servers Tenants								
+ - More Actions								
13 items								
	Name	Size (GB)	Allocated (%)	Used (%)	NAS Server	Pool	Data Reduction	
	File_System_1	200.0			NAS_Server_1	Pool 1	Yes	
	File_System_2	200.0			NAS_Server_1	Pool 1	No	
	File_System_3	1,024.0			NAS_Server_2	Pool 1	Yes	

Figure 12 Unisphere file system space

The figure below shows the total pool space used for this system. This figure also shows the non-base allocated space, which is the amount of pool space used for the snapshot and thin clones, if applicable.

Pool 1 Properties						
General		Drives		Usage	Snapshot Settings	RAID
Capacity		Delete Snapshots		23 items		
Storage Resources						
		Name	Type	Total Pool Space U...	Preallocated (GB)	Non-base Allocated...
<input type="checkbox"/>	✓	File_System_1	File System	20.8	3.1	< 0.1
<input type="checkbox"/>	✓	File_System_2	File System	150.2	2.9	< 0.1
<input type="checkbox"/>	✓	File_System_3	File System	78.5	4.1	0.0
<input type="checkbox"/>	✓	File_System_4	File System	774.7	3.0	0.0
<input type="checkbox"/>	✓	File_System_5	File System	104.4	3.7	0.0
<input type="checkbox"/>	✓	File_System_6	File System	322.5	2.9	0.0
<input type="checkbox"/>	✓	File_System_13	File System	154.0	0.0	0.0
<input type="checkbox"/>	✓	NAS_Server_1	NAS Server	2.2	0.4	0.0
<input type="checkbox"/>	✓	NAS_Server_2	NAS Server	2.2	0.4	0.0
<input type="checkbox"/>	✓	NAS_Server_3	NAS Server	2.2	0.4	0.0
<input type="checkbox"/>	✓	NAS_Server_4	NAS Server	2.2	0.4	0.0
<input type="checkbox"/>	✓	NFSDatastore1	VMware NFS	105.1	6.8	18.0
<input type="checkbox"/>	✓	NFSDatastore2	VMware NFS	576.9	4.7	0.0
<input type="checkbox"/>	✓	VMFS Datastore 1	VMware VMFS	63.1	3.3	0.0
<input type="checkbox"/>	✓	VMFS Datastore 2	VMware VMFS	56.2	6.2	0.0

Figure 13 Unisphere total pool space used

With an understanding of these different values and how they apply to file systems on Dell EMC Unity, we will take a look at the various extend and shrink operations that can be performed on file systems and how each operation affects file system space.

4 Shrink and extend

Dell EMC Unity file systems are built to meet administrators' changing needs as easily and flexibly as possible. Dell EMC Unity allows for increased flexibility by providing the ability to shrink and extend all file system types. With manual and automatic file system extension and shrink with reclaim, Dell EMC Unity makes the most efficient use of pool capacity at all times and allows administrators to respond to changing environmental factors including file system utilization, pool utilization, and client capacity demands. Each of these space efficiency operations can be executed or monitored easily through Unisphere, without requiring administrators to meticulously plan file system size changes or perform complex migrations as requirements change. These operations are also fully compatible with replication. Whenever the source file system is manually or automatically shrunk or extended, the replication destination file system is modified to reflect the same total and allocated space after the next sync completes.

It is important to understand the differences between the manual and automatic shrink and extend operations. Manual shrink and extend operations are used to resize the file system and update the capacity that is seen by the client. This is done by updating the Size attribute in the properties of a file system. If a manual shrink operation on a thin file system shrinks into Allocated space, it may be possible to reclaim capacity back to the pool. For manual shrink and extend, the minimum value is equal to the Used size of the file system and the maximum value is 256 TB. You cannot shrink the file system to less than the Used size, as this would cause the client to see the file system with as 100% full.

Automatic shrink and extend operations occur in the background on thin file systems and do not need to be managed by the administrator. Automatic shrink is designed to ensure file systems are efficiently using their allocated capacity. It continuously checks the used-to-allocated ratio on each file system to ensure the amount of allocated capacity is appropriate. If it detects there is too much capacity allocated to the file system that is not being used, automatic shrink triggers a reclaim of that capacity back to the pool so that it can be used for other resources.

Automatic extend is designed to allocate additional capacity to the file system as it is being written to. When a thin file system is first created, very little capacity is allocated to it regardless of its configured Size. As clients write data to the file system, automatic extend allocates additional capacity from the pool to the file system. This ensures there is enough capacity on the file system to absorb incoming writes. Automatic extend operations continue happening until the Size of the file system is reached.

4.1 Manual extension

When manually extending file systems, only the Size is changed. The allocated and used space remains the same as before the extension. In Unisphere this can be performed by simply changing the Size attribute of the file system from the file system properties page. After extension, the additional space will be visible to the clients of the file system.

4.2 Manual shrink

When an administrator wishes to reduce the client visible file system size and potentially reclaim space to the underlying storage pool, a manual shrink operation can be initiated. This is done in the same way as a manual extension, by changing the file system size in Unisphere to the new desired size. After the shrink operation completes, clients will see the new advertised file system size.

Manual shrink operations may also return unused space to the storage pool, depending on the size of the shrink and the current allocation of the file system. Manual shrink operations can only return space to the pool

if the file system is shrunk into allocated space and will return to the pool a maximum of the difference between the allocated space and the new Size after shrinking. To illustrate this, consider this file system: size = 3 TB, allocated = 1 TB, used = 500 GB. Manually shrinking this file system from 3 TB to 1 TB would return no space to the pool, as the file system was not shrunk into allocated space. However, shrinking from 3 TB to 0.9 TB could potentially return up to 0.1 TB to the pool, depending on the existence of snapshots.

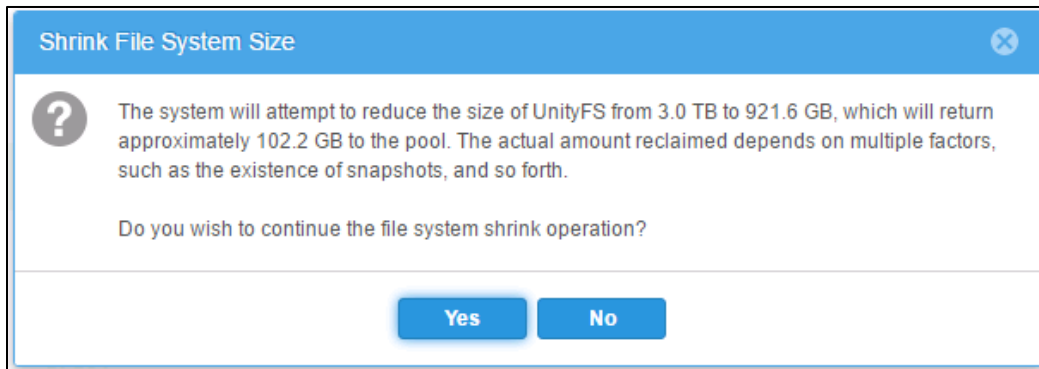


Figure 14 Manual shrink confirmation

It is never possible to shrink into used space, so attempting to reduce the file system size to less than 500 GB would fail in this example. The figure below shows the confirmation message when attempting to shrink a file system in Unisphere, which will calculate the expected amount of space to be reclaimed to the storage pool depending on the current allocation of the file system and requested size. This message indicates that the amount reclaimed depends on the existence of snapshots. This is because snapshots of the file system are required to preserve a view of the file system associated with a particular point in time, and therefore the system cannot allow any blocks associated with an existing snapshot to be reclaimed to the storage pool, even if that space is unused at the current point in time.

For example, suppose a snapshot is taken of a fully allocated 100 GB file system. After taking the snapshot, the administrator immediately shrinks the file system to 80 GB. Because the snapshot must preserve the point-in-time view of the 100 GB file system, the file system shrink operation succeeds in reducing the size of the current production file system, but does not return any space to the storage pool. Despite being shrunk from the production file system, the 20 GB is still associated with the snapshot taken previously, and therefore must be preserved in the event the snapshot needs to be restored in the future. In this circumstance, the confirmation message shows only a very small amount of metadata space to be reclaimed.

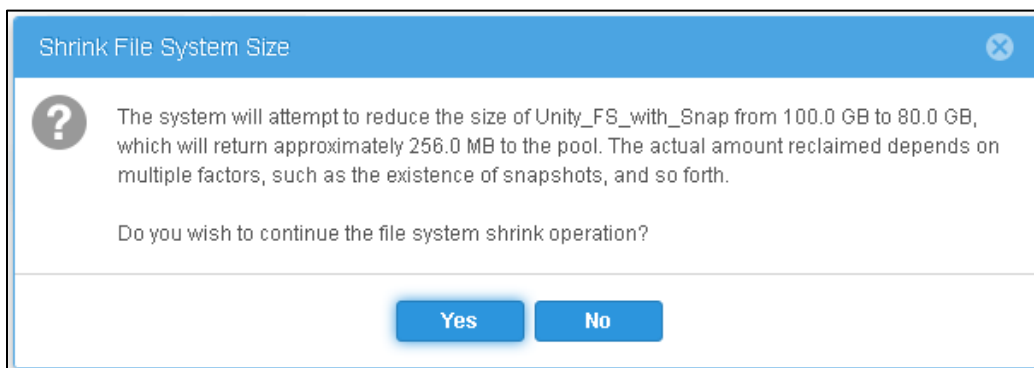


Figure 15 Manual shrink confirmation with snapshot

This is an extreme example used to illustrate the potential effect of snapshots on file system shrink operations, rather than a typical case. The amount of space reclaimed is a function of the amount of changed

data since the last snapshot was taken. The less data that has changed since the last snapshot was taken, the larger the apparent disparity in shrunk data and reclaimed space. In the example above, notice that the shrink operation was initiated **immediately** after taking the snapshot, meaning the snapshot and production file system contained exactly the same data at the time of the shrink operation. Because all 20 GB being shrunk was also tied to the snapshot, none of this space could be returned to the pool. However, if the shrink operation were instead initiated later, after some changes had been made to the data, the snapshot and current file system would no longer be identical. In this case the file system would contain blocks not also associated with any snapshot, and therefore eligible to be reclaimed as part of the shrink operation. For more information on snapshots, reference the *Dell EMC Unity: Snapshots and Thin Clones* white paper on [Dell EMC Support](#).

4.3 Automatic shrink

Under normal operation, Dell EMC Unity file systems automatically adjust the allocated space to optimize storage pool usage. An automatic shrink is this automatic adjustment depending on the ratio of used-to-allocated space. This is because a low used-to-allocated ratio does not represent ideal space utilization, since the allocated but unused space is essentially wasted, and could potentially be used by other pool resources if it were reclaimed to the pool. On Dell EMC Unity, file systems become eligible to be automatically shrunk (have their allocated space reduced) after maintaining an unacceptably low used-to-allocated space ratio for a predefined period of time.

Consider the file system from the previous example: size = 3 TB, allocated = 1 TB, used = 500 GB. The used space is very low relative to the allocated space, utilizing only half of the space reserved from the storage pool. Unless clients begin using additional space from the file system, the Dell EMC Unity system will eventually de-allocate a portion of the 1 TB allocated space back to the underlying storage pool to potentially be used by other resources requiring this space. Automatic shrink only affects the allocated size, so the Size of 3 TB does not change.

In Dell EMC Unity OE version 4.4 and earlier, a file system is eligible for automatic shrink if the used-to-allocated ratio remains under the low watermark, which is 70%. For example, a file system with 400 GB allocated requires 120 GB to be freed before automatic shrink is triggered. Since automatic shrink is based on the used-to-allocated ratio as a percentage, large file systems require more data to be deleted in order to be eligible for automatic shrink. Due to this, file systems that have 400 GB or more allocated use a different mechanism for automatic shrink. For these large file systems, if the gap between the used and allocated space is 20 GB or more, then the file system is eligible for automatic shrink. However, this was not appropriate for very large file systems with greater capacity fluctuations.

In Dell EMC Unity OE version 4.5, the system dynamically adjusts the auto-shrink low watermark depending on the allocated capacity on the file system. This provides a better balance between prematurely shrinking and reclaiming space when appropriate. The chart below shows how the low watermark scales with the file system size.

For example, the low watermark for a file system with 100 TB allocated is 99.089%, as opposed to a low watermark of 70% or a 20 GB gap between the used and allocated space.




Figure 16 Dynamic low watermark

The system also monitors the file system utilization for a period of time prior to initiating an automatic shrink. This avoids prematurely shrinking a file system, only to re-allocate the space when clients begin to write additional data. The system monitors the shrinkable space of the file system every 1.5 hours. It compares the amount of shrinkable space with the results of the previous check and updates a counter. If the amount of shrinkable space is:

- Less than half of the previous check, this indicates the usage has increased significantly
 - Action: Reset the counter back to 0
- More than half of the previous check, this indicates the usage has increased slightly
 - Action: Decrement the counter by 1
- The same or has grown compared to the previous check, this indicates usage is staying the same or decreasing
 - Action: Increment the counter by 1

Once the counter reaches 5, an automatic shrink is initiated. As discussed in the previous section, the actual space reclaimed to the storage pool as a result of a shrink operation will vary based on the existence of snapshots.

4.4 Automatic extension

As the Used space in a file system increases due to more data being written to the file system, more space must be reserved from the storage pool in order to accommodate this new data. As a result, the file system reserves additional space from the pool, increasing the allocated space in the file system. This happens without user intervention and will continue up to the advertised size of the file system. It is important to monitor the amount of free capacity on the pool when overprovisioning thin file systems. If the pool runs out of space, no additional capacity can be allocated to the file system and the file system becomes read-only, until this condition is cleared.

In Dell EMC Unity OE version 4.4 and earlier, the automatic extension is triggered once the used-to-allocated space ratio reaches the high watermark, which is 75%. However, this does not scale well since bigger file systems have large amounts of capacity allocated to it that will not be immediately used. For example, a file system with 100 TB used has an extra 25 TB is allocated from the automatic extend operation. In Dell EMC Unity OE version 4.5, the system dynamically adjusts the auto-extend high watermark depending on the used capacity on the file system. The chart below shows how the low watermark scales with the file system size. For example, the high watermark for a file system with 100 TB allocated is 99.24%, as opposed to a fixed high watermark of 75%.




Figure 17 Dynamic high watermark

Using this method, there is less wasted capacity since the available capacity is used more efficiently. With this enhancement, the same file system with 100 TB used only extends by 761 GB. The graph below shows the extend size based on the used capacity of the file system using the dynamic high watermark functionality. The numbers in the chart show the extend size for file systems with 50, 100, 150, 200, and 250 TB used.




Figure 18 Dynamic extend size

Automatic extension only affects the allocated size, so the size does not change. When the allocated space reaches the size of the file system, no new allocations are made unless an administrator first manually extends the size of the file system.

4.5 Minimum allocation size

On Dell EMC Unity OE version 4.1, file systems can be configured with a minimum allocation size. This enables the storage administrator to control the guaranteed space allocated to a file system. When the minimum allocation size is configured during the creation of a new file system, the size specified is allocated immediately at creation. For example, you could create a 500 GB file system with a minimum allocation size of 100 GB. The host sees a 500 GB file system, but only 100 GB has been allocated to this file system from the pool. Since the remaining capacity is thinly provisioned, additional capacity is allocated as clients consume more than 100 GB of the file system. The minimum allocation size can be configured between 3 GB (default) or up to the size of the file system. This setting can be configured during file system creation and it can be changed at any time.

The minimum allocation size is also used to prevent the automatic shrink feature from reclaiming too much capacity from the file system. For example, if a file system only has 50 GB used but has 100 GB allocated due to the minimum allocation size, automatic shrink would normally be initiated. However, since the minimum allocation size is designed to guarantee allocation to the file system, automatic shrink is not initiated in this scenario. On the other hand, if the minimum allocation size is set to 30 GB but the file system still has 100 GB allocated, automatic shrink is initiated as normal.

Minimum allocation size is compatible with features such as replication and snapshots. When used on a replicated file system, both the source and destination file system have the same minimum allocation size. This is because this setting is propagated to the destination as part of the synchronizations. If a snapshot has a different minimum allocation size setting compared to the file system, restoring the snapshot also restores the minimum allocation size setting.

Starting with Dell EMC Unity OE version 4.2, newly created file systems do not have the option to configure the minimum allocation size. This is due to architectural changes made on the file system to enable support for data reduction. However, Unisphere provides the ability to create thick file systems if space reservation is required. There is no way to convert from thin to thick or thick to thin after a file system is created.

5 File-Level Retention

Dell EMC Unity OE version 4.5 introduces File-Level Retention (FLR). FLR enables the ability to lock files, preventing them from being modified or deleted until a specified retention date. This functionality is also known as write once, read many (WORM). FLR is available on the physical Dell EMC Unity family as well as Dell EMC UnityVSA systems. This feature is only available for file systems and is not available for VMware NFS datastores.

There are two versions of FLR available: Enterprise (FLR-E) and Compliance (FLR-C). FLR-E prevents file modification and deletion by users through access protocols such as SMB, NFS, and FTP. However, an authorized storage administrator can delete the entire file system even if it contains locked files. FLR-C prevents administrators from deleting a file system that contains locked files. The administrator must wait until all files to expire before the file system can be deleted. FLR-C also has other differences including a data integrity check, hard infinite retention, and snapshot restrictions. FLR-C is designed to meet the requirements of SEC Rule 17a-4(f).

Files in a FLR-enabled file system can be in one of the following states:

- **Not locked:** This is the initial state of a file. This file can be modified, moved, renamed, deleted, and so on without any restrictions.
- **Locked:** A locked file can never be modified and can only be deleted once the retention date passes. Files can either be manually locked by users or automatically locked by the system. Once a file is locked, the retention date can be extended but not shortened.
- **Append-only:** Data in append-only files cannot be modified or deleted. However, new data can be added to the end of the file. This functionality is useful for log files. Append-only files can also be locked later.
- **Expired:** A file that was previously locked, but the retention date has passed. Even though it has expired, the data still cannot be modified. Expired files that contain data can be either re-locked or deleted. Empty expired files can be changed to append-only.

For more information on File-Level Retention, reference the *Dell EMC Unity: File-Level Retention (FLR)* white paper on [Dell EMC Support](#).

6 Quotas

Dell EMC Unity includes full quota support to allow administrators to place limits on the amount of space that can be consumed from a user of a file system or directory, or a directory itself in order to regulate storage consumption. These simple but flexible quotas are supported on SMB, NFS, and multiprotocol file systems and can easily be configured through any of the available management interfaces. Due to the particular targeted use case of VMware file datastores, quotas are not available for this resource type.

On Dell EMC Unity systems running OE version 4.4 or earlier, usernames are only displayed when querying for individual user quotas. When viewing quota reports that include multiple users, usernames are not displayed. This is because username lookups need to be queried over the network to the directory service and could potentially take a long time, especially if there are many users in the report. When reviewing quota reports, administrators needed to use the UID to identify each user instead.

Starting with Dell EMC Unity OE version 4.5, an enhancement was added to enable displaying usernames in quota reports. Both UNIX and Windows usernames along with the Windows SID are added. This has several benefits including being much more user friendly, reduces complexity, and improves ease of use.

To eliminate the latency associated with querying usernames on-demand, users are now stored locally on the system. When a new user quota is created, the lookup is done, and the username is stored in the database. When a quota report is generated, usernames are populated directly from the database. The usernames and quota usage in the database are refreshed automatically every 24 hours. If desired, an administrator can initiate an on-demand update of the database by running **uemcli /quota/user {-fs <value> | -fsName <value>} [-path <value>] refresh [-updateNames]**. If the **-updateNames** switch is omitted, then the quote usage is refreshed but, usernames are not. The **-updateNames** option is only available through UEMCLI. Clicking the Reload User Quotas button in Unisphere does not initiate a username database update.

The figure below shows a screenshot that includes the UNIX and Windows Usernames and Windows SID in a quota report.

FS Properties

?

✕

General

Snapshots

FAST VP

Replication

FLR

Quota

Advanced

File System

Quota Tree

User Quota Report

Manage Quota Settings

+

↺

✎

2 items

▼

⚙

▼

<div>☐</div>	<div>...</div>	<div>↑</div>	User ID	Unix N...	Windows Name	Windows SID	Usage (GB)	Soft Limit (...)	Soft Limit Us...	Hard Limit (...)
<div>☑</div>	OK		10000	chen	PRODCSE\Chen	S-1-5-15-5...	0.0	1.0	<div></div>	2.0

Figure 19 Usernames in quota reports

6.1 Quota types

Dell EMC Unity supports file system user quotas, quota trees, and quota tree user quotas. All three types of quotas can coexist on the same file system and may be used in conjunction to achieve finer grained control over storage usage.

File system user quotas are set at a file system level and limit the amount of space a particular user may use from a file system. Administrators also have the ability to choose whether to enforce user quotas for the file system. If quotas are not enforced, they will still be tracked for the file system, but users will not have their file system usage restricted in accordance with the quotas. By default, quotas are not enforced, however this can be changed in the Manage Quota Settings dialog box along with the default user quotas. Default file system level quota limits are applied automatically to all users who access a file system, however these can be overridden for specific users as necessary by creating a new user quota in Unisphere. Because all unspecified users are subject to the default quota settings by default, there is no ability to “delete” user quotas. Instead a user quota can be set to 0 to allow unlimited access, or reset to the default limits, in which case the particular entry would be removed from the user quota list in Unisphere but remain in effect with the default settings.

Quota trees limit the maximum size of a particular directory in a file system. Unlike user quotas, which are applied and tracked on a user by user basis, quota trees are applied to directories within the file system. On Dell EMC Unity, quota trees can be applied on new or existing directories.

If an administrator specifies a nonexistent directory when configuring a new quota tree, the directory will be automatically created as part of quota configuration. However, an administrator may also specify an existing file system directory with existing data when creating a quota tree, allowing the ability to implement quotas on existing file system and directory structures after they have already been in production. If a tree quota is deleted, the directory itself remains intact and all files continue to be available.

Quota trees may not be nested within a single directory. For example, if a quota tree is created on /directory1, another quota tree cannot be created on /directory1/subdirectory1. However, it is possible to have quota trees on /directory2, /directory3, and so on. Once a quota tree has been created, it is also possible to create additional user quotas within that specific directory. Similar to file system level user quotas, the administrator has the option to whether to enforce user quotas specific to this directory and set the default quota limits for the directory. As an example, user Chuck may have a user quota of 25 GB at the file system level, and then the additional restriction of a 10 GB user quota within /directory1 which is limited to 100 GB for all users combined. The figure below shows an illustration of the quota hierarchy possible with the combination of user quotas and quota trees, including directory-specific user quotas.

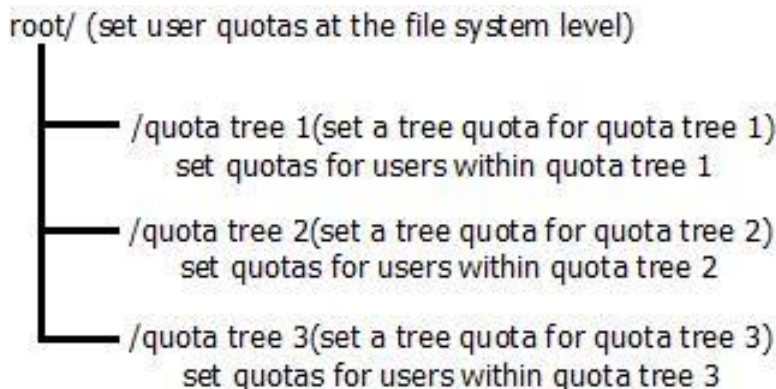


Figure 20 Quota hierarchy

6.2 Quota limits

All quotas consist of three major parameters which determine the amount of space that may be used in a file system in a certain scenario, and define the behavior of the file system or directory when a limit is being approached or exceeded. These parameters are:

- Soft limit (GB)
- Grace period (time)
- Hard limit (GB)

Each of these is configured during quota creation or inherited from the default settings for all quotas. The soft limit is a capacity threshold above which a countdown timer will begin. While the soft limit may be exceeded, this timer, or grace period, will continue to count down as long as the soft limit is exceeded. If the soft limit remains exceeded long enough for the grace period to expire, no new data may be added to the particular directory or by the particular user associated with the quota. However, if sufficient data is removed from the file system or directory to reduce the utilization below the soft limit before the grace period expires, access will be allowed to continue as usual. A hard limit is also set for each quota configured. Upon reaching a hard limit, no new data will be able to be added to the file system or directory. When this happens, the quota must be increased or data must be removed from the file system before additional data can be added. The figure below illustrates this quota behavior.

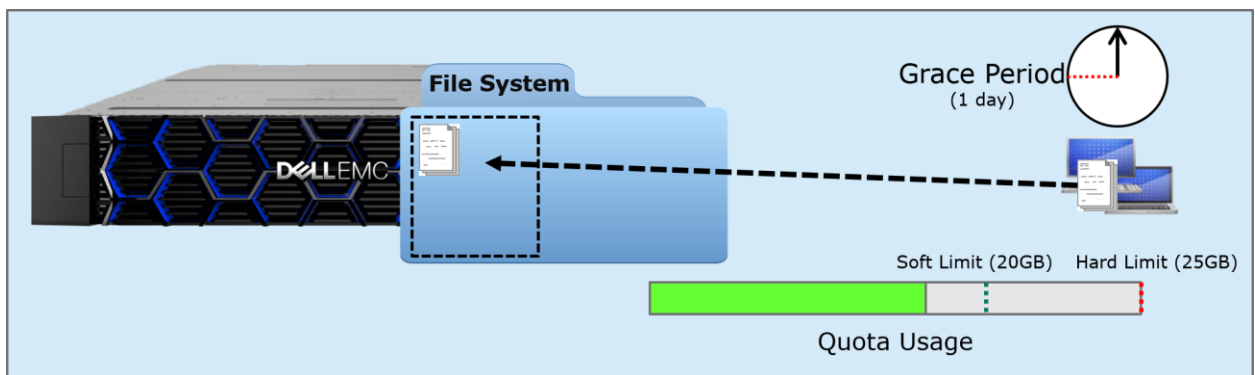


Figure 21 Quota normal operation

Suppose the following user quota has been configured on a file system for a particular user: soft limit = 20 GB, grace period = 1 day, hard limit = 25 GB. The user begins copying data to the file system, and after some time the user has stored 16 GB of files on the file system. Because this is below the limits for the user's quota, the user is still able to add more data to the file system unimpeded.

After some time the user continues to add data to the file system, crossing the 20 GB soft limit. At this point the user is still able to add additional data to the file system, however the grace period of 1 day begins to count down. The storage administrator receives an alert in Unisphere stating that the soft quota for this user has been crossed. If the user does not remove data from the file system prior to the expiration of the grace period, they will no longer be able to add data to the file system until enough data is removed from the file system for the usage to fall below the soft limit.

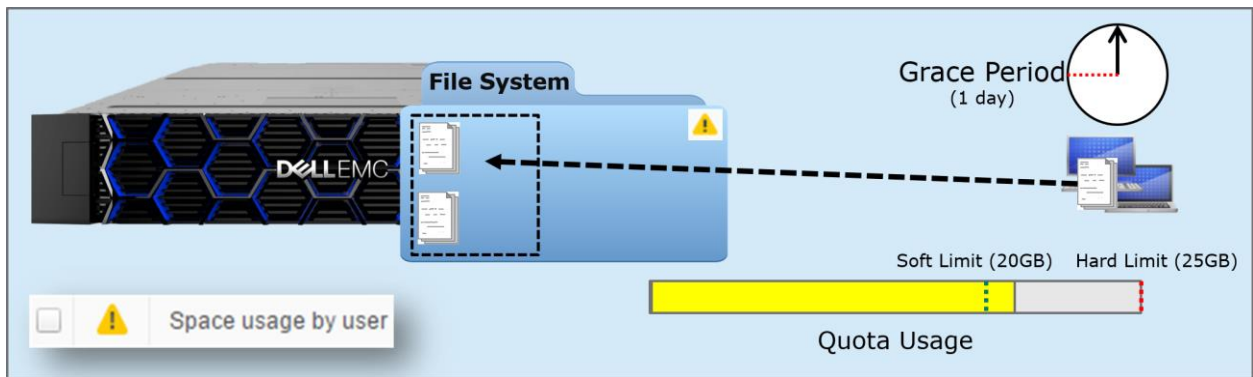


Figure 22 Quota soft limit passed

However, if the user continues writing to and using additional space from the file system despite passing the soft limit, they may eventually reach the hard limit. When this happens, the user will no longer be able to add data to the file system. Administrators will also receive a warning in Unisphere informing them that the hard limit has been reached.

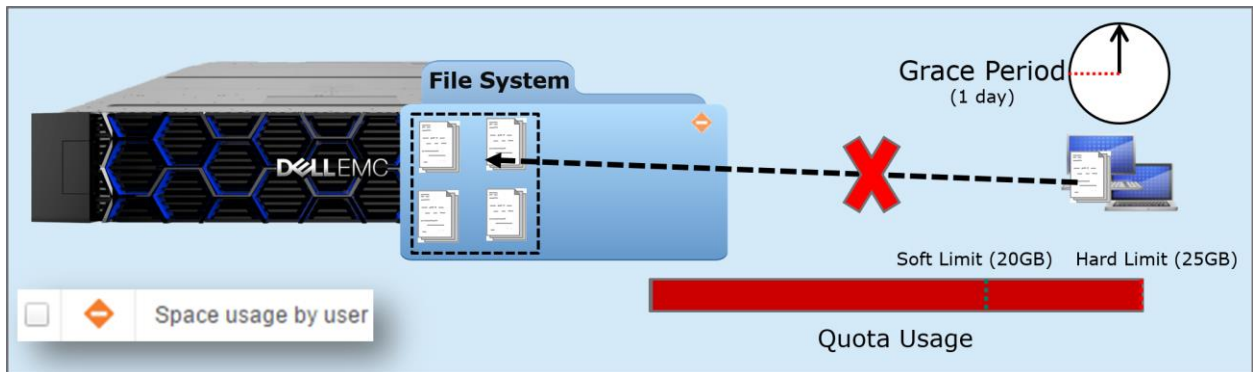


Figure 23 Quota hard limit passed

6.3 Quota policy

When using quotas, administrators have the option to calculate file system usage in one of two ways. This option, which is configured on a per file system basis, may be set to file size or block based calculation. When using the default setting of file size, disk usage is calculated based on logical file sizes in 1K increments. Because of this, it is possible that used space may be reported from a quota perspective as more than the actual usage if holes exist in a sparse file. This setting is generally recommended for Windows environments. The Block quota policy calculates disk usage in 8-16 KB file system blocks, and is accurate with regard to the actual allocation down to the block level. This setting is recommended for UNIX environments.

It is possible to change the quota policy of a file system with existing quotas online, which will initiate a recalculation of the space used for all quotas. If a manual quota recalculation is desired, one can be performed by changing the quota policy and then resetting the policy back to the original setting.

7 Protocol options

Dell EMC Unity supports the concurrent use of all major NAS protocols, including SMB, NFS, FTP and SFTP. Protocol support is configured at the NAS server level, which allows the creation of file systems that can be accessed over that protocol. It is possible for each NAS server to be configured to support one or many different protocols depending on the specific needs of the environment. When enabling a protocol for a NAS server, there are also additional options that can be enabled both at the NAS server and share level.

7.1 SMB

All Dell EMC Unity releases support SMB1 through SMB 3.02, which supports enhancements such as continuous availability, offload copy, protocol encryption, multichannel, and shared VHDX support. Some of these features do not require any special configuration on the Dell EMC Unity system, such as multichannel and shared VHDX support. For multichannel, if there are multiple interfaces created on multiple ports, the SMB3 protocol automatically uses all available TCP connections for a single SMB session. Shared VHDX support provides the ability to enable virtual hard disk sharing on Microsoft Hyper-V® to share a virtual disk between multiple nodes. The configurable options of these features are described later in this section.

Starting with Dell EMC Unity OE version 4.2, SMB 3.1.1 is also supported, which adds reliability enhancements for Continuous Availability (CA) for Hyper-V Cluster Client Failover (CCF) and improved security and encryption traffic performance. The SMB version that is used is dependent on the client operating system.

- **CIFS:** Windows NT 4.0
- **SMB1:** Windows 2000, Windows XP, Windows Server 2003, and Windows Server 2003 R2
- **SMB2:** Windows Vista (SP1 or later) and Windows Server 2008
- **SMB2.1:** Windows 7 and Windows Server 2008 R2
- **SMB3.0:** Windows 8 and Windows Server 2012
- **SMB3.02:** Windows 8.1 and Windows Server 2012 R2
- **SMB3.1.1:** Windows 10 and Windows Server 2016

Regardless of the protocol used for client communication, SMB1 must also be enabled in your environment if running Dell EMC Unity OE version 4.1 or earlier. SMB1 is used to establish a secure channel for communication between the Dell EMC Unity NAS server and the domain controllers in your environment. This secure channel is used for operations such as authentication, SID lookups, Group Policies, and so on. If SMB1 is disabled in your environment, these operations will fail. However, starting with Dell EMC Unity OE version 4.2, SMB1 no longer needs to be enabled in your environment since SMB2 is used for secure channel communication, by default. Using SMB2 enhances security and increases efficiency due to enhancements and updates to the protocol. This allows customers that have security concerns or company policies to disable SMB1 without impacting connectivity. If SMB2 is not available, the Dell EMC Unity NAS server attempts to use SMB1 as a backup option. This means that any domain controllers that are running older operating systems that only support SMB1 can continue to function. In addition, you can also disable SMB1 for client access on the NAS server by using the **cifs.smb1.disabled** parameter. For more information about NAS server parameters and how to configure them, reference the *Service Commands* document on [Dell EMC Support](#).

SMB support is enabled on the NAS server level during or after creation, allowing administrators to create SMB-enabled file systems on that NAS server. When enabling SMB support on a NAS server, the SMB server can either be standalone or Active Directory domain joined. Domain joined NAS servers are placed in the OU=Computers, OU=EMC NAS servers organizational unit, by default.

Dell EMC Unity also supports the Microsoft Distributed File System (DFS) namespace. This provides the administrator the ability to present shares from multiple file systems through a single mapped share. A Dell EMC Unity SMB server can be configured as a standalone DFS root node or as a leaf node on an Active Directory DFS root. DFS-R (replication) is not supported on Dell EMC Unity systems. If replication is required, the native asynchronous replication feature can be used to replicate the file system instead.

Each SMB file system and share has additional advanced protocol options that are disabled by default but can be set by administrators. SMB protocol related options are shown in the table below.

Table 2 SMB options

Protocol options	Level	Default
Sync Writes Enabled	File system	Disabled
Oplocks Enabled	File system	Enabled
Notify On Write Enabled	File system	Disabled
Notify On Access Enabled	File system	Disabled
Continuous Availability	Share	Disabled
Protocol Encryption	Share	Disabled
Access-Based Enumeration	Share	Disabled
Branch Cache Enabled	Share	Disabled
Offline Availability	Share	None
UMASK (Multiprotocol Only)	Share	022

7.1.1 Sync writes enabled

Synchronous writes enable the storage system to perform immediate synchronous writes for storage operations, regardless of how the SMB protocol performs write operations. Enabling synchronous writes operations allow you to store and access database files (for example, MySQL) on storage system SMB shares. This option guarantees that any write to the share is done synchronously and reduces the chances of data loss or file corruption in various failure scenarios, for example, loss of power. If SMB3 Continuous Availability (CA) is enabled, all write operations are automatically synced to satisfy the requirements for CA. This option can have a big impact on performance. It is not recommended unless you intend to use Windows file systems to provide storage for database applications.

7.1.2 Oplocks enabled

Opportunistic file locks (oplocks) allow SMB clients to buffer file data locally before sending it to a server. SMB clients can then work with files locally and periodically communicate changes to the storage system rather than having to communicate every operation over the network to the storage system. Unless your

application handles critical data or has specific requirements that make this mode or operation unfeasible, leaving the oplocks enabled is recommended.

The following oplocks implementations are supported on Dell EMC Unity:

- **Level II Oplocks:** Informs a client that multiple clients are currently accessing a file, but no client has yet modified it. A level II oplock lets the client perform read operations and file attribute fetches by using cached or read-ahead local information. All other file access requests must be sent to the server.
- **Exclusive Oplocks (SMB2 only):** Informs a client that it is the only client opening the file. An exclusive oplock lets a client perform all file operations by using cached or read-ahead information until it closes the file, at which time the server must be updated with any changes made to the state of the file (contents and attributes).
- **Batch Oplocks:** Informs a client that it is the only client opening the file. A batch oplock lets a client perform all file operations by using cached or read-ahead information (including opens and closes). The server can keep a file opened for a client even though the local process on the client machine has closed the file. This mechanism curtails the amount of network traffic by letting clients skip the extraneous close and open requests.

This option only applies to client access over SMB1 since oplocks are always enabled for client access over SMB2. However, disabling this option also invalidates the SMB2.1 file and directory lease feature. Leasing serves the same purpose as oplocks, but provides greater flexibility and enhancements, increasing performance and reducing network utilization.

- **Read-caching lease:** Allows caching reads and can be shared by multiple clients.
- **Write-caching lease:** Allows caching writes and is exclusive to only one client.
- **Handle-caching lease:** Allows caching handles and can be shared by multiple clients.

7.1.3 Notify on write/access enabled

This option enables notifications when a file system is written to or accessed. Applications that run on Windows platforms, and use the Win32 API, can register with the SMB server to be notified of file and directory content changes, such as file creation, modify, or rename. For example, this feature can indicate when a display needs to be refreshed (Windows Explorer) or when the cache needs to be refreshed (Microsoft Internet Information Server), without having to constantly poll the SMB server.

7.1.4 Continuous availability

Continuous availability is an SMB3+ specific feature that can be enabled at the share level on Dell EMC Unity systems. In the event of a client or storage processor failure, CA allows persistent access to Dell EMC Unity file systems without loss of the session state. This is useful for critical applications such as Hyper-V or SQL, where constant availability to files is of the utmost importance. SMB 3.0 uses persistent handles to enable the Dell EMC Unity NAS server to save on disk-specific metadata associated to an open handle. In the event of an SP failure, applications accessing open file content are not affected as long as the NAS server and file system failover to the peer SP completes within the timeout of the application. This results in clients transparently reconnecting to the peer SP after the NAS server failover without affecting those clients' access to their files.

Continuous availability is also available on the client side, which is independent from storage CA. Client CA transparently preserves access in the event of a node failure within a client application cluster. When a failure of one node in the cluster occurs, the application is moved to the other node and reopens its content on the

share from that node using its originally assigned **ApplicationID** without an interruption in access. The CA option on the share does not need to be enabled in order to utilize client CA.

Starting with Dell EMC Unity OE version 4.2, SMB 3.1.1 is supported. This adds a reliability enhancement for Continuous Availability for Hyper-V Cluster Client Failover by adding an **ApplicationInstanceVersion** tag in addition to the **ApplicationID**. The **ApplicationInstanceVersion** tag is incremented each time an application is restarted on a new node within the cluster. In situations where network access is lost, but storage access remains available, the application may be restarted on a new node without the cluster knowing due to the lack of network access. The **ApplicationInstanceVersion** tag enables the storage system to easily identify which node in the cluster is the correct owner of the application. The storage system can safely close any locks that were opened with a lower **ApplicationInstanceVersion** number, which allows the application to restart without any conflicts.

7.1.5 Protocol encryption

Protocol encryption is an SMB 3.0 feature that is available on Dell EMC Unity. This option provides in-flight data encryption between SMB 3.0 compatible clients and the Dell EMC Unity NAS server. Data is encrypted by the client before being sent to the NAS server, or from NAS server to client. It is then decrypted upon reaching its destination, whether that is the NAS server or SMB client. The protocol encryption is enforced at user session level, ensuring the whole SMB traffic is encrypted once the user session is established.

The following setting can be configured in the NAS server's registry:

- **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\LanmanServer\Parameters\RejectUnencryptedAccess:** Determines if clients that do not support encryption (pre-SMB3.0) have access to the share
 - 1 (default): Returns access denied to pre-SMB3.0 clients that do not support encryption
 - 0: Allows pre-SMB3.0 clients to access the share without encryption

Starting with Dell EMC Unity OE version 4.2, SMB 3.1.1 is supported. This provides improved security and encryption traffic performance for SMB3 by changing the encryption algorithm from AES-CCM-128 to AES-GCM-128. This change improves performance under certain conditions such as large file transfers. In addition, this improves security against man-in-the-middle attacks.

7.1.6 Access-based enumeration

Access-based enumeration is a share-level option that restricts the display of files and folders based on the access privileges of the user attempting to view them. Without access-based enumeration, all users are able to view all files and folders within a directory to which they have access. However they will not be able to open or view these files and folders without the appropriate access privileges. When access-based enumeration is enabled on a share, users will only be able to see files or folders for which they have at read access or above. For example, without access-based enumeration a user without access to several files would still be able to see that those files exist in a directory to which they have access. However with access-based enumeration, that same user would not even see those same inaccessible files in the directory. Administrator users are always able to see all files and folders, even when access-based enumeration is enabled on a share.

7.1.7 BranchCache

BranchCache is a share-level option that allows users to access data stored on a remote NAS server locally over the LAN without being required to traverse the WAN to access the NAS server. This is most useful in a remote or branch office environment, where branch offices are required to access data stored on a remote

server at the main office. BranchCache allows this data to be cached locally at the branch, either by a single designated BranchCache server or distributed across clients, in order to reduce WAN bandwidth used by many clients constantly and repeatedly traversing the WAN for the same data.

With BranchCache enabled, the client uses the WAN to retrieve the hash of the file from the remote NAS server. The searches the local file cache to look for a file with a matching hash. If all or some of the data is available locally, either on the designated BranchCache or another client computer, the data will be retrieved locally. The data is validated using a hash function to ensure the file is exactly the same. Any data that is not cached locally is retrieved from the NAS server over the WAN, and then cached locally for future requests. BranchCache works best for data that does not change often, allowing files to be cached for longer periods of time at the branch offices.

7.1.8 Offline availability

Offline availability is a share-level attribute that allows administrators to determine if and how files and programs in a share will be available when offline. This allows users to access shares on a server even when they are not connected to the network by storing a version of the share in a local cache on the client computer. In order for offline availability to function, it must be configured on both the share and the individual client computers accessing the share. Dell EMC Unity NAS servers support four options for offline availability, which are the same options supported by Windows file servers and are shown below.

- **Manual (Default):** Only files and programs that the users specify will be available offline. Nothing will be cached without the user requesting it.
- **Cache all programs and files opened by users:** All files and programs that users open from the share are automatically available offline. Whenever a user accesses a file or program from a share, that content will automatically be cached so as to be available to that user in offline mode. All files opened will continue to be cached and available for offline access until the cache becomes full or the user deletes particular files from the cache. Cached content will continue to sync with the version on the server. Files and programs that have not been opened will not be available offline.
- **Cache all programs and files opened by users, optimize for performance:** The same as above, except that executable files that have been previously cached locally will be run from the cached copy rather than the copy on the share, even when the share is available. This option is useful for reducing network traffic and performance overhead.
- **None:** No files or programs from the share will be available offline. Client computers will not be able to cache any content from this share for offline access.

7.2 NFS

Dell EMC Unity supports NFSv3 through NFSv4.1. NFS can be enabled on a per NAS server basis, which only enables the NFSv3 protocol. After enabling NFSv3, administrators have the option to enable additional options such as vVols, NFSv4, and Secure NFS. Afterward, when creating an NFS file system and share, default host access can also be set on a share level with exceptions defined for specific hosts.

In Dell EMC Unity OE version 4.4 and earlier, NFSv3 must be enabled before NFSv4 can be enabled. Dell EMC Unity OE version 4.5 adds the ability to enable NFSv3 or NFSv4 independently. This is useful for customers who only use NFSv4 and want to leave the NFSv3 protocol disabled. NFSv4 is not supported with vVols.

Starting with Dell EMC Unity OE version 4.4, NFS share names can contain the "/" character, except as the first character. Previously, using the "/" character in the share name is prohibited as it is reserved to indicate a directory on UNIX systems. By allowing the use of the "/" character in the share name, this enables

administrators to create a virtual name space that is different from the actual path used by the share. While NFSv3 clients can mount shares that include “/” in the name, NFSv4 client cannot due to protocol constraints. NFSv4 clients must use the actual path of the share to mount it.

By default, NFS clients have the ability to set the **setuid** and **setgid** bits on files and directories stored on an NFS export. Files that the **setuid** or **setgid** bits set can be identified since the execute bit changes from x to s for the owner or group, respectively. If set on a file, these bits provide the ability for users to run an executable with the permissions of the owner or group, respectively. Also, if the **setgid** bit is set on a directory, it causes new files and subdirectories to inherit the GID from its parent directory, rather than the primary GID of the user that is creating the file. The **setuid** bit is ignored if it is set on a directory on most UNIX and Linux systems. These bits allow users to run specific executables with temporarily elevated permissions (such as root) which may be a security concern for some customers. Dell EMC Unity OE version 4.4 also introduces the ability to allow or prevent clients from setting the **setuid** and **setgid** bits on any files and directories residing on the NFS share. By default, this is allowed and can be changed when creating or modifying an NFS share. If it is disabled on an existing NFS share, any existing files that already have the **setuid** or **setgid** bits configured are not changed. However, any future attempts to set or unset these bits are not allowed. Also, if files that have these bits enabled are copied to the share, these bits are removed. The allow SUID option is shown in the figure below.

FS Properties

General | Host Access

Share Name: FS

Description:

NAS Server: NAS1

File System: FS

Local Path: /FS/

Export Paths: 10.10.10.10:/FS

☒ Allow SUID

Anonymous UID: 4294967294

Anonymous GID: 4294967294

Close Apply

Figure 24 Allow SUID and anonymous UID/GID

Dell EMC Unity OE version 4.4 also introduces the ability to configure the anonymous UID and GID attributes. If a client is granted access to a NFS share without allowing root access, the root user on that client may still attempt to access the share. In this case, the root user is mapped to anonymous UID and GID 4294967294, which is typically associated with the **nobody** user. A custom anonymous UID and GID can be configured on the NFS share during creation or modification of the share as shown in the figure above. By default, these are

set to 4294967294 and can be changed to any valid UID and GID. If the client is granted root access to the share or if Secure NFS is used, the anonymous UID and GID are not used.

NFS protocol-related options are shown in the table below.

Table 3 NFS options

Protocol options	Level	Default
NFSv4	NAS server	Disabled
Secure NFS (with Kerberos™)	NAS server	Disabled
vVols (NFS Protocol Endpoint)	NAS server	Disabled
Default Host Access	Share	No Access
Allow SUID	Share	Enabled
Anonymous UID and GID	Share	4294967294

7.2.1 Parameters

By default, NAS servers query lookup services in the following order: local files, LDAP/NIS, and then DNS. In Dell EMC Unity OE version 4.5, the **ns.switch** parameter is available. This parameter is used to control which services are queried and also the order that they are searched. It has the ability to specify different services for different types of object lookups. The available options are **passwd**, **group**, **hosts**, and **netgroups**. The available resolvers are **files**, **nis**, **ldap**, and **dns**. The default is NULL, which means the default order is used. If you want to use a custom lookup order, edit this using the same syntax as the **nsswitch.conf** file, but concatenate the lines. For example, if you only want to use local files for user lookups and only DNS for host lookups, set this to **passwd: files hosts: dns**.

When using a string to define the host access list for an NFS share, netgroups should be prefixed with an @ symbol to differentiate them from hostnames. When using registered hosts, the @ symbol is not required since these entries are defined according to their host object type.

In the string, the system treats all entries that begin with @ solely as netgroups. Entries that do not begin with @ are resolved as hostname first and, if that fails, as a netgroup second. This could cause a performance impact due to unnecessary and unintentional netgroup lookups. Dell EMC Unity OE version 4.4 and later adds a parameter to disable the netgroup lookup behavior for unresolved hostnames. The **nfs.netgroupprefix** parameter can be configured to:

- 0 (default)
 - Entries that begin with @: Treated only as a netgroup
 - Entries that do not begin with @: Treated as a hostname first and as a netgroup second
- 1
 - Entries that begin with @: Treated only as a netgroup (no change)
 - Entries that do not begin with @: Treated only as a hostname

The NFSv4 protocol allows the owner and group attributes of a file to be identified as either a numeric ID or a string in the packet. Some legacy NFSv4 clients require the use of the string format as they do not support numeric IDs. Dell EMC Unity OE version 4.4 adds two parameters to control this behavior, enabling support for some of the legacy clients. Parameters **nfsv4.numericid** and **nfsv4.domain** are described below.

- **nfsv4.numericid**: Specifies if NFSv4 user and group attributes are handled as numeric IDs or user@domain strings
 - 0: Uses strings (user@dell.com)
 - 1 (default): Uses numeric IDs (UID 100 and GID 100)

If **nfsv4.numericid** is set to 0, a UNIX Directory Service (UDS) must be configured in order to translate the numeric UID and GID to the usernames. The domain is retrieved from a second parameter:

- **nfsv4.domain**: Specifies the domain to be used for the usernames. This is only used if **nfsv4.numericid** is set to 0. If this parameter is empty, the realm of the NFS server (if Secure NFS is enabled) is used as the domain name.

NFSv4 supports file delegations which is the process of assigning management of a file to the client. This greatly reduces the number of transactions required between the NAS server and the client which results in increased efficiency, reduced network traffic, and potentially improved performance. Read delegations can be provided to multiple clients simultaneously since they only prevent write access while the data is being read. Write delegations are exclusive and are used to prevent read access by other clients while new data is being written. Since NFSv4 delegations provide several advantages, they are enabled by default. However, in specific instances such as troubleshooting or issue reproduction, this feature may need to be disabled. Disabling this feature may result in increased network traffic. Dell EMC Unity OE version 4.3 also includes a parameter to disable NFSv4 delegation. The **nfsv4.delegationsEnabled** parameter can be configured to:

- 0: Disables NFSv4 file delegations
- 1 (default): Enables NFSv4 file delegations

Dell EMC Unity systems use NFS export aliases, meaning each NFS export has a path and also a name alias. An NFS client can mount the NFS export by using either the path or name. For example, if the path is **/filesystem1/** and the export name is **FS1**, either can be used to mount the export. In addition, both options also show up in the **showmount -e** output even though they are the same export. In order to prevent duplicates in the **showmount** output, a parameter is available in Dell EMC Unity OE version 4.3 to control this behavior. The **nfs.showExportLevel** parameter can be configured to:

- 0 (default): Show both the export path and name
- 1: Show only the export path
- 2: Show only the export name

NFS is designed as a simple and efficient protocol. Oracle takes this a step further with their Oracle® direct NFS (dNFS) client. dNFS is built into the database's kernel, enabling performance improvements compared to traditional NFS. A traditional NFS kernel goes through four layers of the I/O stack while the dNFS only needs to go through two. This is accomplished by bypassing the operating system level caches and eliminating operating system write-ordering locks. This enables dNFS to generate precise requests without any user configuration or tuning. Memory consumption is also reduced since data only needs to be cached in user space, eliminating the need for a copy to exist in the kernel space. Performance can also be further enhanced by configuring multiple network interfaces for load balancing purposes. All Dell EMC Unity OE versions support Oracle dNFS in single node configurations. Starting with Dell EMC Unity OE version 4.2,

Oracle Real Application Clusters (RAC) are also supported. In order to use Oracle RAC, the **nfs.transChecksum** parameter must be enabled. This parameter ensures that each transaction carries a unique ID and avoids the possibility of conflicting IDs that result from the reuse of relinquished ports. The **nfs.transChecksum** parameter can be configured to:

- 0 (default): Does not compute a CRC
- 1: Computes a CRC on the first 200 bytes to check if an XID entry matches the request

For more information about NAS server parameters and how to configure them, reference the *Service Commands* document on [Dell EMC Support](#).

7.2.2 NFSv4

NFSv4 is a version of the NFS protocol that differs considerably from previous implementations. Unlike NFSv3, this version is a stateful protocol, meaning that it maintains a session state and does not treat each request as an independent transaction without the need for additional preexisting information. This behavior is similar to that seen in Windows environments with SMB. NFSv4 brings support for several new features including NFS ACLs that expand on the existing mode-bit-based access control in previous versions of the protocol.

While Dell EMC Unity fully supports the majority of the NFSv4 and v4.1 functionality described in the relevant RFCs, directory delegation and pNFS are not supported. Some of these features do not require any special configuration on the Dell EMC Unity system, such as NFSv4.1 session trunking. In order to enable this, create multiple interfaces on the NAS server and then point the host to all of the available IP addresses.

To configure NFSv4, you must first enable NFSv4 on the NAS server and then create an NFS file system and share. Then, the file system can be mounted on the host using the NFSv4 mount option.

Starting with Dell EMC Unity OE version 4.2, NFS datastores can also be mounted using NFSv4. When creating NFS datastores on earlier versions of Dell EMC Unity OE, the NFSv3 protocol is always used. If you want to use NFSv4, ensure NFSv4 is enabled on the NAS server. When creating a new datastore, select NFSv4 on the configuring host access page and provide host access to the ESXi servers. This process creates the datastore and automatically mounts it on the ESXi servers using NFSv4.

7.2.3 Secure NFS

Traditionally, NFS is not the most secure protocol, because it trusts the client to authenticate users as well as build user credentials and send these in clear text over the network. With the introduction of secure NFS, Kerberos can be used to secure data transmissions through user authentication as well as data signing through encryption. Kerberos is a well-known, strong authentication protocol where a single key distribution center, or KDC, is trusted rather than each individual client. There are three different modes available:

- **krb5**: Use Kerberos for authentication only
- **krb5i**: Use Kerberos for authentication and include a hash to ensure data integrity
- **krb5p**: Use Kerberos for authentication, include a hash, and encrypt the data in-flight

In order to enable secure NFS, both DNS and NTP must be configured. Also, a Unix Directory Service such as NIS, LDAP, or Local File must be enabled, and a Kerberos realm must exist. LDAPS (LDAP over SSL) is generally used for Secure NFS to avoid weaknesses in the security chain. If an Active Directory domain joined SMB server existed on the NAS server, that Kerberos realm may be leveraged. Otherwise, a custom realm can be configured for use in Unisphere.

Starting with Dell EMC Unity OE version 4.2, NFS datastores can also be mounted using Secure NFS with Kerberos. When creating NFS datastores on earlier versions of Dell EMC Unity OE, the NFSv3 protocol is always used. If you want to use Secure NFS with Kerberos, ensure Secure NFS is enabled on the NAS server. You must also configure DNS, NTP, domain, and NFS Kerberos credentials on the ESXi server. When creating a new datastore, select NFSv4, provide the Kerberos NFS Owner name, and provide either read/write or read-only access to your ESXi hosts. This process creates the datastore and automatically mounts it to the ESXi hosts using secure NFS.

7.2.4 vVols

Virtual Volumes (vVols) is a storage framework introduced in VMware vSphere 6.0 that is based on the VASA 2.0 protocol. vVols enable VM-granular features and Storage Policy Based Management (SPBM). Enabling this option allows clients to access NFS vVol datastores through a NAS server.

Using vVols with IP multi-tenancy is not supported. You are prohibited from enabling the vVol Protocol Endpoint on a NAS server that has a tenant association. In order to use vVols, you must use a NAS server that has NFSv3 enabled and does not have a tenant assigned.

For more information on vVols, reference the *Dell EMC Unity: Virtualization Technology* white paper on [Dell EMC Support](#).

7.2.5 Host access

The default host access option determines the access permissions for all hosts with network connectivity to the NFS storage resource. The available options are:

- No Access (Default)
- Read-Only
- Read-Only, allow Root (Dell EMC Unity OE version 4.4 or later for file systems)
- Read/Write
- Read/Write, allow Root

For hosts that need something other than the default, different access levels can be configured by adding hosts, subnets, or netgroups to the override list with one of the access options above. To do this, these resources must first be registered on to the system. This can be done in the **Hosts (ACCESS)** page in Unisphere. The following information is required for registration and all other fields are optional:

- Host: Name and IP Address/Hostname
- Subnet: Name, IP Address, and Subnet Mask/Prefix Length
- Netgroup: Name and Netgroup Name

For hostname resolution, the search order is local files, UNIX directory service, and then DNS. This means if Local Files are configured, they are always queried first to resolve a hostname. If the name cannot be resolved or if Local Files are not configured, then the NAS server queries the configured UNIX Directory Service (LDAP or NIS), if it is configured. If the name still cannot be resolved or if neither Local Files nor UNIX Directory Service is configured, then DNS is used. This behavior can be customized by configuring the **ns.switch** parameter on the NAS server.

Starting with Dell EMC Unity OE version 4.4, NFS host registration is made optional. Instead, host access can be managed by specifying a comma separated string. This is designed to simplify management and improve ease of use. When configuring access to a NFS share, you can select the option to enter in a comma

separated list of hosts or select from a list of registered hosts on the system. For each NFS share, only one host access method can be used.

As part of this change, you may see the host registration-based method of configuring NFS access referenced to as advanced host management, such as in UEMCLI. If advanced host management is enabled, it indicates the NFS share is using the host registration-based method for configuring NFS access. If Advanced Host Management is disabled, it indicates the NFS share is using the comma separated string method of configuring NFS access.

The string can contain any combination of entries listed in the table below and is limited to 7000 characters. If replication is configured, this string is also replicated to the destination, so no reconfiguration of host access is required in the event of a failover.

Table 4 NFS host access

Name	Example	Notes
Hostname	host1.dell.com	Hostname should be defined in the local hosts file, NIS, LDAP, or DNS.
IPv4 or IPV6 Address	10.10.10.10 fd00:c6:a8:1::1	
Subnet	10.10.10.10/255.255.255.0 10.10.10.10/24	IP address/netmask or IP address/prefix.
Netgroup	@netgroup	Netgroup should be defined in the local netgroup file or UDS. Netgroup entries should be prefixed with @ to differentiate them from hostnames.
DNS Domain	*.dell.com	The DNS Server must support reverse lookups and the ns.switch parameter should not exclude DNS. Domain entries should be prefixed with * and follow the Linux convention. This option is only available when using host strings.

7.3 Multiprotocol

When configuring a NAS server for protocol access, an administrator has several options. With respect to SMB and NFS, the NAS server can be configured in one of the following ways:

- SMB only
- NFS only
- SMB and NFS (separate SMB and NFS file systems)
- Multiprotocol (SMB and NFS to the same file system)

The major difference between enabling both SMB and NFS independently and enabling multiprotocol is that multiprotocol configurations allow data in a single file system to be accessed through both SMB and NFS concurrently. In contrast, a non-multiprotocol NAS server with SMB and NFS enabled individually will require separate file systems to be configured for SMB and NFS, where SMB users will not be able to access NFS

file system data and NFS users cannot access SMB file system data. When a NAS server is designated as multiprotocol, all file systems on that NAS server will be multiprotocol accessible.

Due to the inherent differences between the SMB and NFS protocols, some configuration is required in order to support multiprotocol. For example, Windows uses security identifiers (SIDs) while UNIX uses user IDs/group IDs (UIDs/GIDs) so various services are required to translate usernames to the proper IDs. Another example is security where Windows uses access control lists (ACLs) while UNIX uses mode bits or the NFSv4 ACL. The following section provides additional details about multiprotocol configuration on Dell EMC Unity.

7.3.1 Directory services

In order to use multiprotocol, you must first configure a multiprotocol NAS server. This provides the configuration that enables the NAS server to recognize users across multiple protocols. In order to enable multiprotocol on a NAS server, the following requirements must be met:

7.3.2 SMB

An Active Directory domain-joined SMB server is required to translate SIDs to Windows names. To configure an Active Directory domain-joined SMB server, NTP must be configured on the system and DNS must be configured on the NAS server. Then, navigate to the **NAS server Properties → Sharing Protocols → SMB** and provide the SMB computer name, Windows domain, domain privileged username, and password.

7.3.3 NFS

At least one of the services below must be configured in order to translate UIDs to UNIX names:

- UNIX directory service (UDS): LDAP or NIS
- Local files: Available in Dell EMC Unity OE version 4.1 or later

To configure LDAP, navigate to the **NAS server Properties → Naming Services → LDAP/NIS**. Provide the LDAP servers, Base DN, Authentication method, and credentials (if necessary). On this page, you can also configure the LDAP schema, enable LDAP secure (use SSL) to encrypt LDAP traffic, and configure the Certification Authority (CA) certificate for authentication. You can configure LDAP to use anonymous, simple, or Kerberos authentication.

In Dell EMC Unity OE version 4.3, the ability to run LDAP lookups from the NAS server is available. This is useful for confirming the mappings are configured properly and also for troubleshooting purposes. You can look up a user, group, UID, GID, host, or netgroup. In order to run an LDAP lookup, run the **svc_nas <NAS_Server> -ldap -lookup** command.

In addition, support for dynamic LDAP domain lookups on NAS servers is added. This feature enables the ability to automatically obtain the LDAP server IP addresses and ports from DNS. Dynamic LDAP domain lookups is the default option and removes the requirement for the user to manually specify the LDAP server IP addresses and ports when configuring or editing a NAS server. This also enables the ability to dynamically and globally update the configuration without modifying each individual NAS server. In order for this to work, SRV (service) records for each LDAP server must exist in DNS and all servers should share the same authentication settings.

Any LDAP server IP addresses that are discovered through this method are not displayed in Unisphere. If you want to view the discovered IP addresses and settings, run the **svc_nas <NAS_Server> -ldap** command to display the current configuration. The system automatically refreshes the configuration every 20 minutes from

DNS. You can also manually refresh the configuration at any time by running the **svc_nas <NAS_Server> - ldap --refresh** command.

If the system is replicating to a system that is running Dell EMC Unity OE version 4.2 or earlier, the discovered IP addresses are configured as static IPs on the destination system. Also, note that this feature only resolves the LDAP server IP addresses and ports, and the rest of the configuration still needs to be entered by the administrator. Also, if desired, the administrator still has the ability to configure the LDAP server IP addresses manually. The dynamic LDAP domain lookups feature is displayed in the figure below.

Create a NAS Server

- General
- Interface
- Sharing Protocols
- Unix Directory Service**
 - DNS
 - Replication
 - Summary
 - Results

Configure Unix Directory Service

☐ Use local files

☒ Enable a Unix Directory Service using NIS or LDAP LDAP

☒ Obtain LDAP servers IPs automatically

☐ Configure LDAP servers IPs manually

Authentication: * Simple

☐ LDAP Secure (Use SSL)

Distinguished Name: * CN=Administrator,CN=Users,DC=dellemc,DC=com

Password: * *****

Base DN: * DC=dellemc,DC=com

Profile DN:

For accessing multiprotocol file systems, your Unix Directory Service must provide the Unix UID and GUID for each Active Directory user.

You can customize the mapping of user names between the Active Directory and Unix Directory Service after the NAS server is created.

You can upload local files (password, group, hosts, netgroup, and homedir) from the NAS server details page after the NAS server is created.

Cancel Back Next

Figure 25 Dynamic LDAP domain lookups

The description and syntax for all of the LDAP configuration settings are shown in the table below.

Table 5 LDAP configuration settings

	Anonymous	Simple (iPlanet or OpenLDAP)	Simple (AD LDAP or IDMU)	Kerberos
Server	LDAP Server IPs or Hostnames (Not required when using dynamic LDAP domain lookups)			
Port	LDAP Server Port Number - Default: 389 / SSL: 636 (Not required when using dynamic LDAP domain lookups)			
Base DN	Base DN in LDAP notation format. For example, if using svt.lab.com, the Base DN would be DC=svt,DC=lab,DC=com	The Base DN is the same as the Fully Qualified Domain Name. For example, svt.lab.com	Base DN in LDAP notation format. For example, if using svt.lab.com, the Base DN would be DC=svt,DC=lab,DC=com	
Profile DN (optional)	Profile DN for the iPlanet or OpenLDAP server			
Distinguished name (simple only)		User account in LDAP notation format. For example, cn=administrator,cn=users,dc=svt,dc=lab,dc=com		
Password (simple only)		User account password		
Notes	Active Directory (AD) is not supported with Anonymous LDAP authentication			See below for Kerberos authentication options

There are two methods for configuring Kerberos:

- **Authenticate to the SMB domain:** Authenticate using the SMB server account or authenticate with other credentials.
- **Configure a custom realm:** Point to any type of Kerberos realm (Windows, MIT, or Heimdal). With this option, the NAS server uses the custom Kerberos realm defined in the Kerberos subsection of the NAS server's Security tab. AD authentication is not used when you choose this option. If using NFS secure with a custom realm, you have to upload a keytab file.

The LDAP configuration must adhere to either the IDMU, RFC 2307 or RFC2307bis schemas. Refer to the RFC for a list of what is required for each schema. You can verify the current schema configuration by using the **Retrieve Current Schema** link on the LDAP page to retrieve the **ldap.conf** file, edit it, and upload a new

version. All containers specified in the **ldap.conf** file must point to a location that is valid and exists in Active Directory, including ones that may not be in use, such as netgroup and host. If any entries are removed from this file, the NAS server automatically sets them to a default value based on the baseDN, which may result in lookup issues. Consult with your domain administrator to get the proper values for each container. The figure below shows an example of a valid LDAP schema for IDMU.

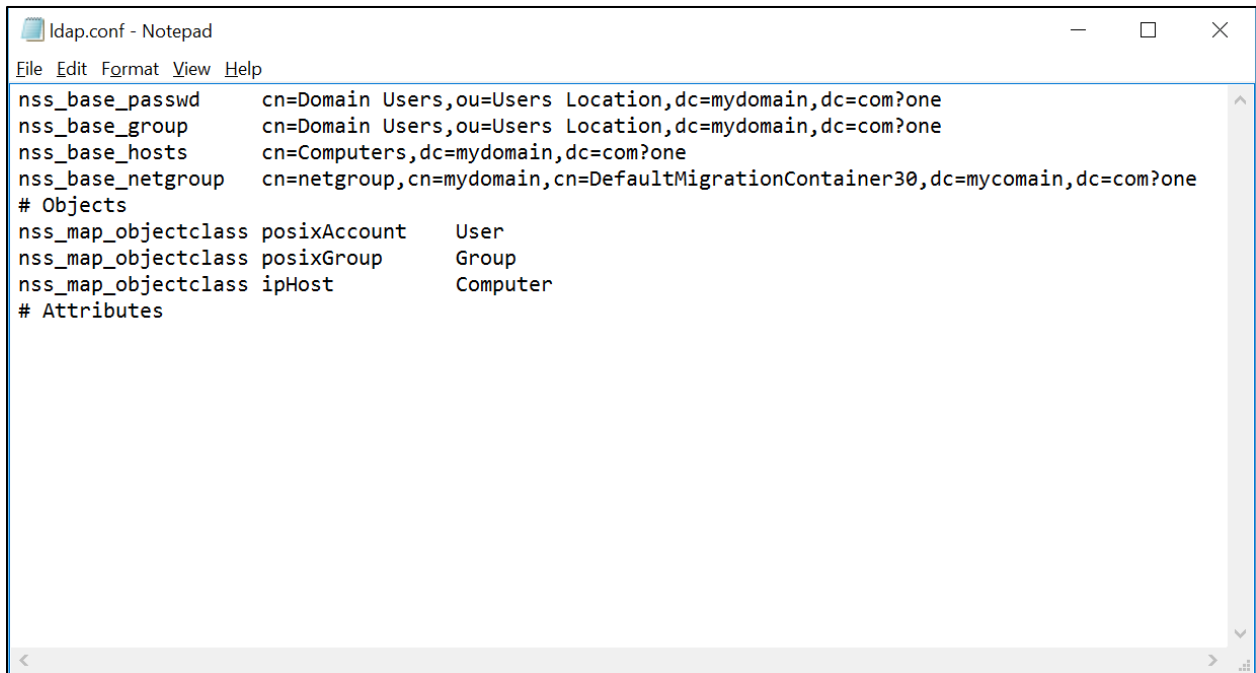


Figure 26 LDAP.conf

To configure NIS, navigate to **NAS server Properties → Naming Services → LDAP/NIS**. Then, provide the NIS domain and up to three NIS servers.

Starting with Dell EMC Unity OE version 4.1, local files can also be used for resolving UNIX identities. This leverages the local passwd and group files that are uploaded to the NAS server. This can be used in addition to or instead of the UDS. If both are configured, Local Files are searched first before using the UDS. To configure Local Files, navigate to **NAS server Properties → Naming Services → Local Files**. You can download the current local files from this page, which also provides syntax and additional details. After the files are edited, upload the new files back to the NAS server. The figure below shows the syntax and an example entry in a **passwd** file. The comment, home directory, and shell should be empty since they are not used.

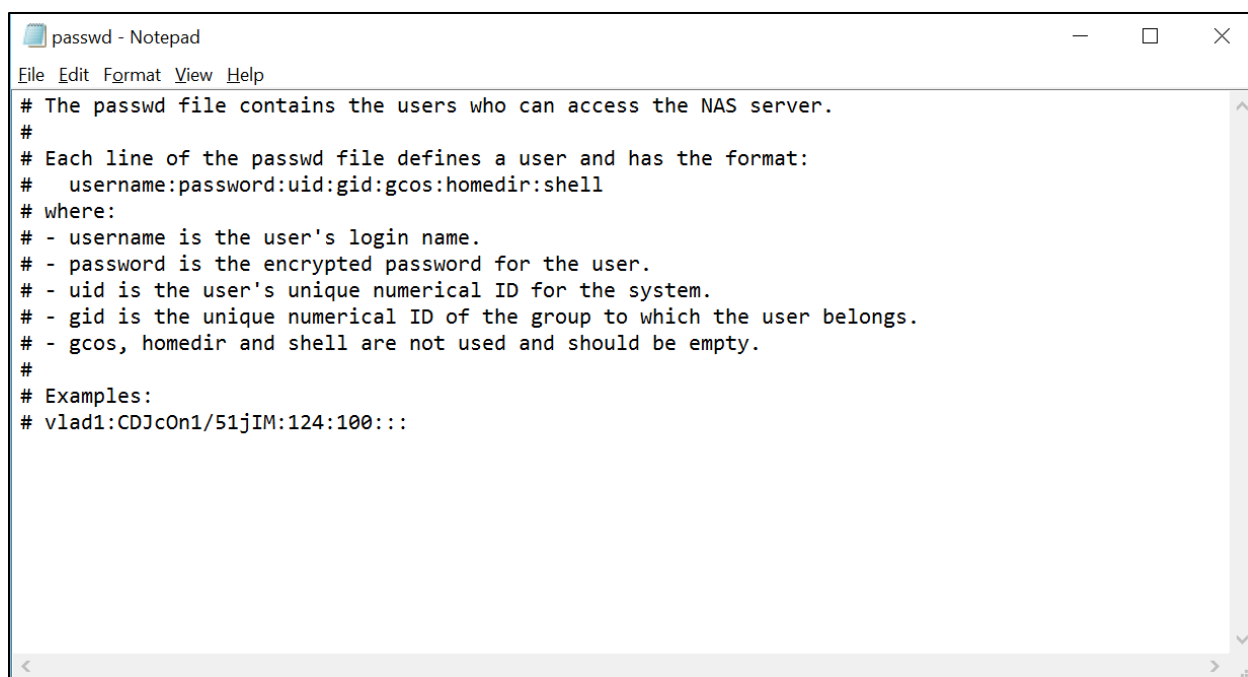


Figure 27 Passwd file

Once all of these requirements are fulfilled, multiprotocol access can be enabled on the NAS server and then a multiprotocol file system can be created. It is important to note that once multiprotocol is enabled and a file system is created, it cannot be disabled. The figure below shows a NAS server that has multiprotocol enabled along with the NTXMAP and mapping diagnostics options.

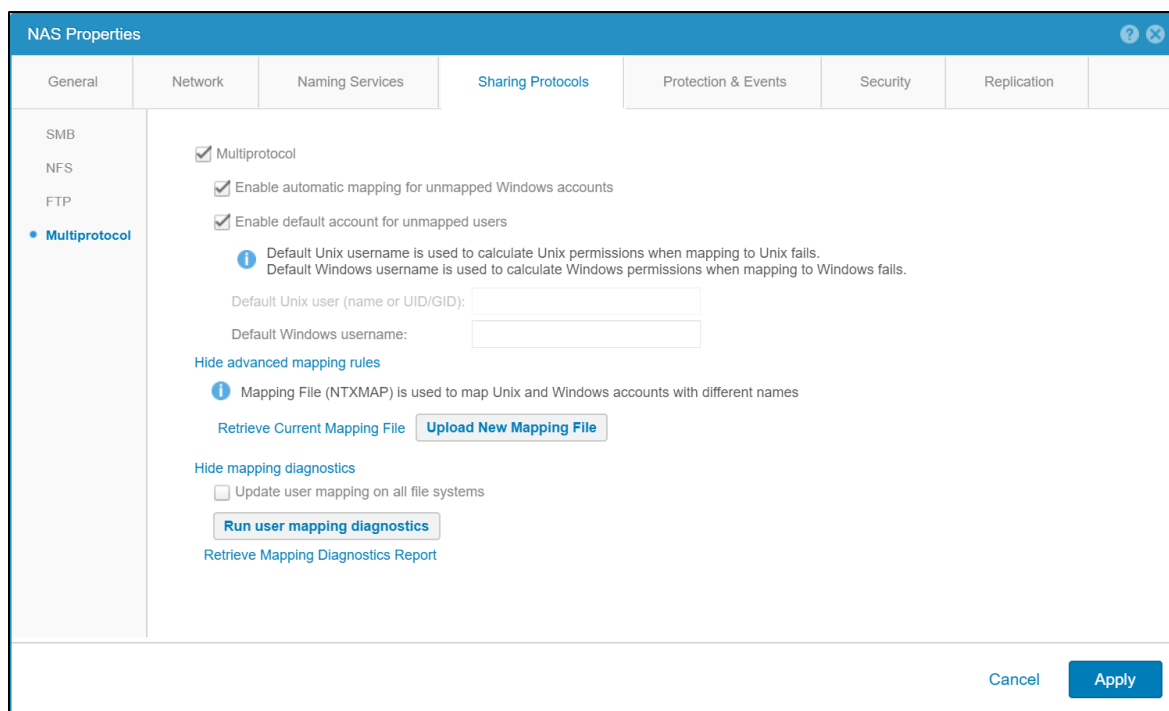


Figure 28 Multiprotocol NAS server

Multiprotocol NAS servers can only support multiprotocol file systems. They do not have the ability to support SMB-only or NFS-only file systems. If multiprotocol is enabled on an existing NAS server, all of the existing file systems are automatically converted to multiprotocol. The file system access policies are automatically updated for UNIX for the existing NFS file systems and Windows for the existing SMB file systems. The system then begins to update the mappings between the Windows and UNIX accounts.

7.3.4 User mapping

Using a multiprotocol file system allows each SMB user to be mapped to a corresponding NFS user with common access privileges, regardless of which protocol the user is using to access their file system data. By default, users with the same name in Active Directory and the UNIX directory service will have their SMB and NFS identities mapped, allowing for multiprotocol access across each protocol. For example, if a user possesses a windows domain user account named **charles** and a UNIX LDAP account also named **charles**, he would be able to access his data with the same privileges across either protocol while being identified as the same user.

However, if his Windows domain user account name was **charles** but his UNIX LDAP account name was **chuck**, his Windows and Linux user names would not be mapped to one another and identified as the same entity by default. If users have different Windows and UNIX usernames, NTXMAP needs to be configured so they can be identified as the same user. It is important to note that NTXMAP is an optional component since it does not provide UID to name mappings. It is a user-defined and managed file that is only used to provide a translation for users that have different names. To configure NTXMAP, navigate to **NAS server Properties → Sharing Protocols → Multiprotocol → Show advanced mapping rules**. From here, you can download the current mapping file, make the appropriate edits, and upload the new file to the NAS server. Once this has been done in our example, Charles is able to access his same data whether accessing the file system through SMB as **charles** or through NFS as **chuck**.

For multiprotocol, the NAS server needs to know the mapping between the SID, SMB Name, UNIX Name, and UID. This provides the ability for a Windows user to be matched to a UNIX user, and conversely, in order to enforce file security when the other protocol is used for access. This cross-protocol mapping is principally done by matching names between the protocols, but each protocol also requires a method to map their respective names to their IDs.

The components involved in user mapping are shown in the table below.

Table 6 User mapping

Name	Service	Description
Windows resolvers (SMB)	Local group database (LGDB) or domain controller	LGDB is used for SMB-only access. DC is used to return: Windows account name for a SID SID for a Windows account name
UNIX Directory Service (NFS)	LDAP/NIS, Local Files, or Both	Used to return: UNIX account name for a UID UID and primary GID for a UNIX account name
Secure Mapping Cache	Secmap Cache	A local cache that maintains all the mappings used by a NAS server to ensure coherency across multiple file systems. The following mappings are tracked: SID to UID and primary GID UID to SID
ntxmap	NTXMAP	NTXMAP is used to associate a Windows account to a UNIX account when the names are different.

7.3.5 Default accounts for unmapped users

Default accounts for unmapped users allow administrators to designate a specific existing Windows and/or UNIX account to serve as the mapping destination for unmapped users wishing to access file system data over the other protocol. For example, in an environment where many users have only Windows accounts, a default UNIX user may be designated in order to allow these unmapped users to access the multiprotocol file system. If default users are not enabled, unmapped users are denied access when attempting to access a multiprotocol file system.

With default accounts enabled, the UID and primary GID of the default UNIX user are used if an unmapped Windows user attempts to access the file system through NFS. Similarly, the credentials of the default Windows user are used when an unmapped UNIX user attempts to access the file system through SMB. The configured default users must be valid and exist in the UDS/local files or DC/LGDB for this to work properly. Although multiple users could write to the file system as the default user, this user is still considered a single user for quota calculation purposes and the UNIX account may have ownership of files from many different Windows users.

On Dell EMC Unity OE version 4.2 and earlier, the default UNIX user can be configured by specifying a username, as long as an entry for that username exists in the UDS/local files to resolve the UID/primary GID. Starting with Dell EMC Unity OE version 4.3, the default UNIX user can be optionally configured as numerical UID/primary GID value. This enables the ability to configure a default UNIX user without setting up and creating a user in the UDS/local files. The specified UID must be in the 32-bit range and follow this format:

@uid=<UID>,gid=<GID>@. For example, if you want to configure a default UNIX user with a UID 1000 and primary GID of 2000, enter **@uid=1000,gid=2000@**. The figure bellows shows how to configure default accounts for unmapped users.

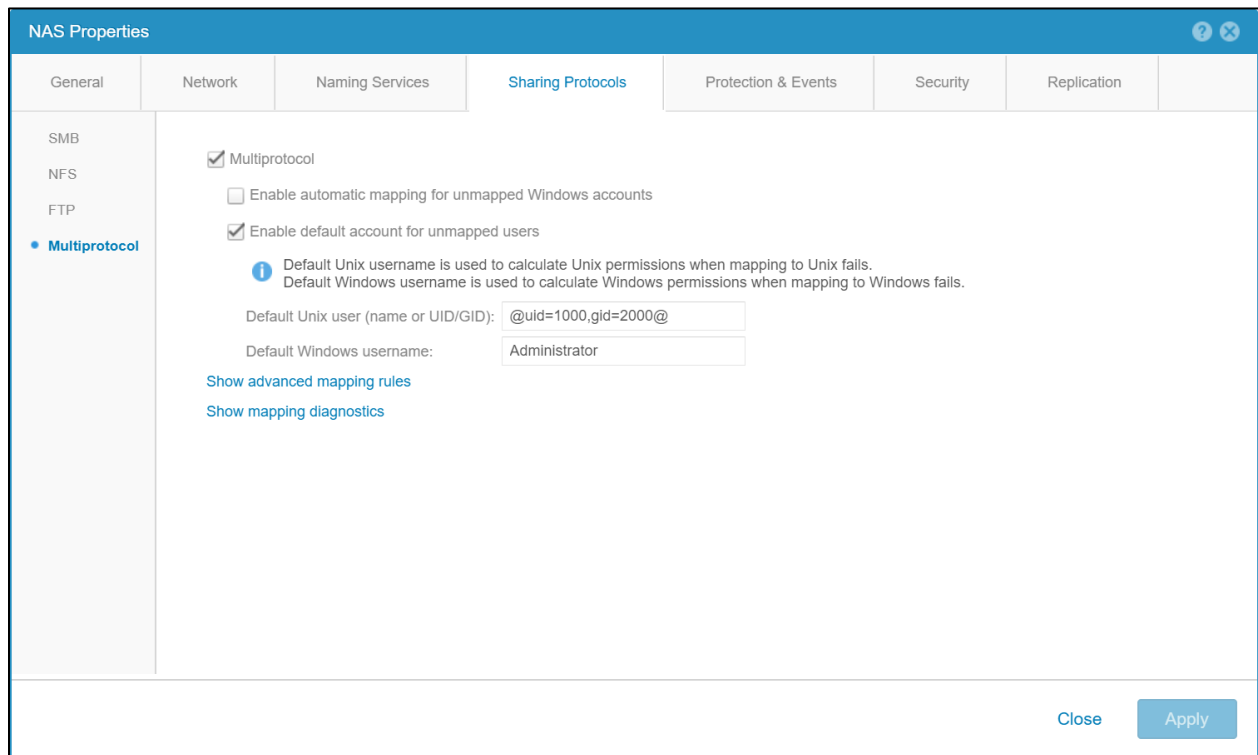


Figure 29 Multiprotocol NAS server

7.3.6 Automatic mapping for unmapped Windows accounts

Starting with Dell EMC Unity OE version 4.3, the option to enable automatic mapping for unmapped Windows accounts is available. This feature enables the ability to automatically generate and assign a unique UID to Windows users that do not have a UID mapping. This feature enables access to the share for unmapped users. Previously, access is denied for unmapped users unless the default UNIX user is configured. Using this feature provides an advantage compared to using the default UNIX user since the unique UIDs enables user quotas to be tracked and enhances security. When using a default UNIX user, multiple users may be writing to the file system using the same default account so user quotas cannot be tracked and the UNIX account may have ownership of files from many different Windows users. If this feature is enabled, the ability to configure the default UNIX username is disabled. This feature is designed for multiprotocol environments that consists of mostly Windows users and the actual UID is not critical.

This feature generates 32-bit UIDs with the most significant bit set to prevent conflicts with UIDs defined by the administrator in the UDS/local files. The range of UIDs generated by this feature is between 2147483649 (0x80000001) and 2151677951 (0x803FFFFF). The automatic UID is only assigned if the user does not already have a UID configured in the UDS/local files.

It is important to note that if the UDS/local files is updated to configure a UID after one is already assigned by this feature, the new entry in the UDS/local files is ignored. If you would like to use the entry in UDS/local files, you must delete the entry from secmap cache by running **svc_cifssupport <NAS_Server> -secmap -delete -name <Name> -domain <Domain>**. Alternatively, you can initiate a full remapping of the entire file system if multiple updates are required.

7.3.7 Mapping process

The figure below shows the process used to resolve a Windows user (SID) to a UNIX user (UID/primary GID). Even when multiprotocol is configured, local users on the SMB Server can still be used for SMB-only access. These users do not need a mapping to a UNIX user as they are only used for SMB access.




Figure 30 SID to UID and primary GID mapping

1. Secmap is searched for the SID. If the SID is found, the UID and primary GID mapping is resolved.
2. If the SID is not in secmap, the Windows username must be found.
 - a. The local group database (LGDB) is searched for the SID to determine if this is a local user. If the SID is found, the name is SMB_SERVER\USER. Since this is a local user for SMB-only access, no UNIX mapping is required.
 - b. If SID is not found in the LGDB, the DC is searched for the SID. If the SID is found in the domain, the name is DOMAIN\USER.
 - c. If the SID is not resolvable, access is denied. This failed mapping added to the persistent secmap database.
3. If the default UNIX account is not used, the Windows name is translated to the UNIX name.
 - a. If the Windows name is found in NTXMAP, that entry is used as the UNIX name.
 - b. If the Windows name is not found in NTXMAP or if NTXMAP is disabled, the Windows name is used as the UNIX name.
4. The local files or UDS are searched for the UNIX name to find the UID and primary GID.
 - a. If the UNIX name is found, the UID and primary GID mapping is resolved. This successful mapping is added to the persistent secmap database.
 - b. If the UNIX name is not found, but the automatic mapping for unmapped Windows accounts feature is enabled, the UID is automatically assigned. This successful mapping is added to the persistent secmap database.
 - c. If the UNIX name is not found, but a default UNIX account is configured, the UID and primary GID are mapped to that of the default UNIX account. This failed mapping added to the persistent secmap database.

- d. If the UNIX name is not resolvable, access is denied. This failed mapping is added to the persistent secmap database.

Mappings that do not result in a permanent UID are considered failed mappings – 2c, 4c, and 4d. Users with failed mappings are added to the secmap database with 4294967294 as their UID. This indicates that the mapping process needs to be retried the next time the user connects. If a mapping is defined for these users at a later time, they can convert in to successfully mapped users upon connecting. The secmap database is then updated accordingly with the permanent mapping. This means these users must go through the mapping process each time they connect until they have a permanent mapping defined.

The figure below shows the process used to resolve a UNIX user (UID) to a Windows user (SID). This process is only needed if the access policy is set to Windows. This is different compared to the UNIX UID that is always required, regardless of the access policy, since the UID is also used for quota management purposes.




Figure 31 UID to SID mapping

1. Secmap is searched for the UID. If the UID is found, the SID mapping is resolved.
2. If the UID not in secmap, the UNIX username must be found.
 - a. Local files and/or UDS are searched for the UID. If the UID is found, the UNIX name is determined.
 - b. If the UID is not found, but a default Windows account is configured, the UID is mapped to the default Windows account. If it does not already exist, the default Windows user is added to the persistent secmap database.
 - c. If the UID is not resolvable, access is denied.
3. If the default Windows account is not used, the UNIX name is translated to the Windows name.
 - a. If the UNIX name is found in NTXMAP, that entry is used as the Windows name.
 - b. If the UNIX name is not found in NTXMAP or if NTXMAP is disabled, the UNIX name is used as the Windows name.
4. The DC or LGDB are searched for the Windows name to find the SID.

- The Windows name is searched in the DC. If the Windows name is found, the SID mapping is resolved. This successful mapping is added to the persistent secmap database.
- If the Windows name contains a period (.) and the part of the name following the last period (.) matches an SMB server name, the LGDB of that SMB Server is searched. If the Windows name is found, the SID mapping is resolved. This successful mapping is added to the persistent secmap database.
- If the Windows name is not found, but a default Windows account is configured, the SID is mapped to that of the default Windows account. If it does not already exist, the default Windows user is added to the persistent secmap database.
- If the Windows name is not resolvable, access is denied.

7.3.8 Mapping management and diagnostics

Starting with Dell EMC Unity OE version 4.3, the ability to manage secmap cache is available. This provides the ability to create, update, and delete any or all entries from secmap cache. This is in addition to the existing options to view, import, export, and generate reports for secmap cache. In order to manage secmap cache, run the **svc_cifssupport <NAS_Server> -secmap** command. There are options to list the contents of secmap cache, create entries, update entries, delete entries, import a database, export the database, or generate a report on the database health and content.

Dell EMC Unity storage also provides the ability to view a user mapping diagnostics report. This report only displays any user mappings that have issues. This process compares the users in the secmap cache to the entries in Active Directory and UDS/local files and generates a report. Under **NAS server Properties → Sharing Protocols → Multiprotocol**, click **Show mapping diagnostics** and **Run user mapping diagnostics**. After the report is generated, select **Retrieve Mapping Diagnostic Report** to download the report. You can view the user mapping issues and make any changes, if necessary. After corrections are made, you can optionally re-run the user mapping diagnostics report to ensure they are no longer on the list.

In the report, each line represents a user that has a mapping issue in one of the formats described in the table below.

Table 7 Mapping diagnostics

Format	Status	Issue
u sid	Unknown user mapping	This is an entry for an unknown SID that cannot be resolved to a username. The SID no longer exists in the DC and may have been deleted.
u sid username	Unknown user mapping	This is an entry that is resolved successfully to a username, but not to a UID. There is no entry in the UDS/local files for this username.
n sid username Uid OLD UID	Known user mapping	This is entry that is resolved successfully to both a SID and UID, but the entry has changed. This means the UID returned from the UDS/local files does not match the UID in secmap cache.

The figure below shows an example of a report with users that have an inconsistent UID and unresolved SID/UID mapping.

```

secmap_report.txt - Notepad
File Edit Format View Help
# User Mapping Diagnostic Files# lines starting by U means: the SID to uid resolution has failed
# lines starting by N means: the resolution has succeeded
# The format of the line for resolved entries is:
#   N <sid> <username> <uid> [<olduid>]
# The format of the line for unresolved entries is:
#   U <sid> [<username>]
#username is present if the SID/name resolution has succeeded but not the name/uid
N S-1-5-15-13a441e3-8c2bf4bb-28a0a9b-1f4 user1 1000 1001
U S-1-5-15-5c9a296e-44073dfc-395c14b-6fd user2
  
```

Figure 32 Mapping diagnostic report

7.3.9 Additional options

When creating or managing a multiprotocol NAS environment, there are additional configuration options at the NAS server and file system levels related to the mapping between SMB and NFS users accessing file system data. These options are shown in the table below.

Table 8 Multiprotocol options

Option	Level	Default
Access policy	File system	Native
UMASK (SMB)	Share	022

7.3.10 Access policy

Security is also handled differently between SMB and NFS. For NFS, the authentication credentials are provided by the client or, for secure NFS, built from the UDS. User rights are determined by the NFSv3 mode bits or NFSv4 ACLs and the UID/GIDs are used for identification purposes. There are no privileges associated with UNIX security.

For SMB, the credentials are built from the domain controller (DC) and local group database (LGDB) of the SMB server. User rights are determined by the ACL and the SID is used for identification purposes. Windows security includes privileges such as **TakeOwnership**, **Backup**, and **Restore** that are granted by the LGDB or group policy object (GPO).

The access policy is used to define how security is enforced on a multiprotocol file system. The default setting of Native maintains two separate sets of permissions for each file and the protocol that is used to access the file determines which set of permissions are checked. If SMB is used, the ACLs are checked but if NFS is used, the NFSv3 mode bits or NFSv4 ACL are checked. If the multiprotocol environment is heavily weighted toward users of one type or another, setting the access policy to one of the other values may be desirable. The Windows setting only checks the ACLs and completely ignores the NFSv3 mode bits and NFSv4 ACL while the UNIX policy does the opposite. The table below describes the available policies that can be configured at the file system level.

Table 9 Access policy

Access policy	Description
Native (default)	<p>Manages access for each protocol separately with its own native security</p> <p>Security for NFS shares use the UNIX mode bits or NFSv4 ACL</p> <p>Security for SMB shares use the SMB Access Control List (ACL)</p> <p>The two sets of permissions are independent and there is no synchronization between them</p> <p>NFSv3 UNIX mode bits or NFSv4 ACL permission changes are synchronized to each other, but SMB ACL is not changed</p> <p>SMB ACL permission changes do not change the NFSv3 UNIX mode bits or NFSv4 ACL</p>
Windows	<p>Uses the SMB ACL for both protocols</p> <p>Upon request for NFS access, the Windows credential built from the DC/LGDB is used to check the ACL for permissions</p> <p>NFSv3 UNIX mode bits or NFSv4 ACLs are updated when SMB ACL permissions are changed</p> <p>NFSv3 UNIX mode bits or NFSv4 ACL permission changes are denied</p>
UNIX	<p>Uses the NFSv3 UNIX mode bits or NFSv4 ACL for both protocols</p> <p>Upon request for SMB access, the UNIX credential built from the UDS/local files is used to check the NFSv3 mode bits or NFSv4 ACL for permissions</p> <p>SMB ACL permissions are updated when NFSv3 UNIX mode bits or NFSv4 ACLs are changed</p> <p>SMB ACL permission changes are allowed to avoid causing disruption, but these permissions are not maintained</p>

The figure below shows how to configure the access policy in Unisphere.

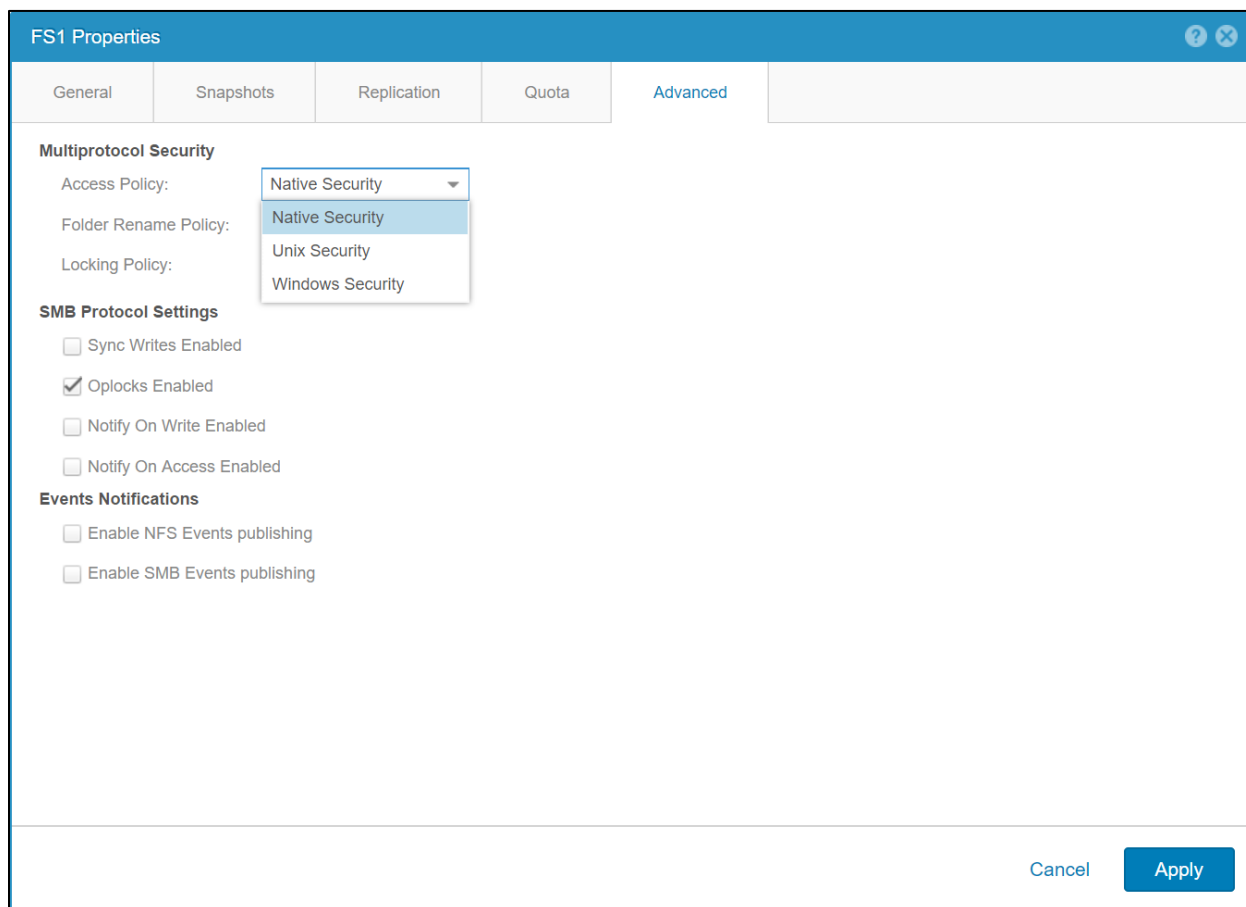


Figure 33 File system access policy

7.3.11 UMASK

The UMASK is a bitmask that enables the ability to control the default UNIX permissions for newly created files and folders. This setting is only applied to new files and folders created on SMB for multiprotocol file systems. The UMASK setting determines which UNIX permissions are excluded for new files and directories. By default, new files have 666 permissions while new directories have 777 permissions. If the UMASK is set to the default value of 022, this means new files have 644 permissions and new directories have 755 permissions instead. If NFSv4 ACL inheritance is present on the directory, this will take precedence over the UMASK setting.

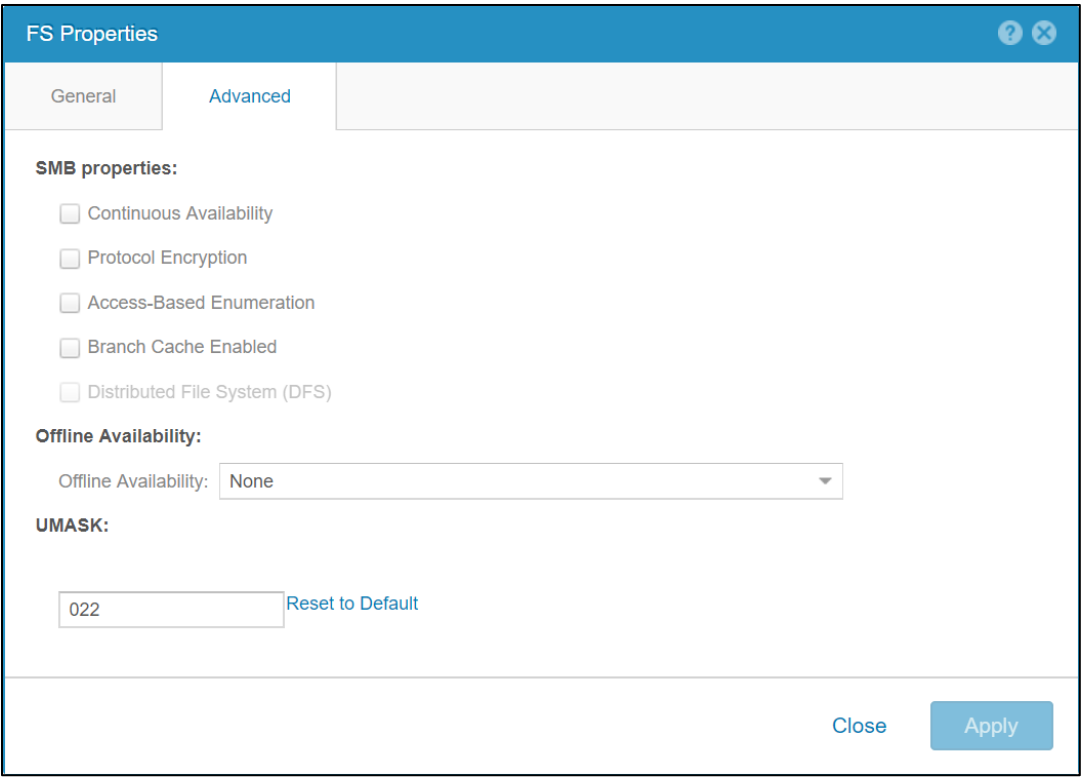


Figure 34 UMASK

On Dell EMC Unity OE version 4.2 and earlier, the UMASK setting is not used for the Windows access policy. Instead, the NAS server performs a conversion of the Windows permissions to determine the proper UNIX permissions. Although the UNIX permissions are not used while the Windows access policy is active, these permissions are still generated. This provides the ability to change the access policy without needing to re-permission the entire file system. The way the permissions are calculated are shown in the table below.

Table 10 Permissions conversion

User	Description
Owner	The rights of the owner and the rights of everyone from the ACL are used.
Group	A logical OR is done on the entries that are not exactly the owner in the ACL to determine the permissions.
Other	Receives the same permissions as the primary group.

Starting with Dell EMC Unity OE version 4.3, this method is no longer used to determine the UNIX permissions. Instead, the default permissions for files created using the other protocol are always determined by the UMASK setting or ACL inheritance. For files created on SMB, the NFSv3 UNIX permissions are determined by the UMASK setting. For files created on NFS, the SMB ACLs are determined by ACL inheritance. This behavior is now consistent regardless of the access policy that is configured.

This behavior is only used to determine the UNIX permissions when creating new files. If permissions are changed on an existing file, the behavior depends on the configured access policy.

For more information about configuring and troubleshooting Multiprotocol, reference the *Configuring Multiprotocol File Sharing and Service Commands* documents on [Dell EMC Support](#).

7.4 Locking and folder rename policy

Starting with Dell EMC Unity OE version 4.1, the locking and folder rename policies can be configured on multiprotocol file systems. These settings allow the administrator to control the desired behavior since locking and folder renaming behave differently depending on the protocol. Both of these settings can be configured during file system creation or afterwards.

7.4.1 Locking policy

Range locks allow hosts to lock a byte range of a file. These locks can be shared locks (writes denied) or exclusive locks (reads/writes denied). Each protocol leverages either mandatory or advisory locking. For mandatory locks, any IO to the locked range is denied. For advisory locks, it is the client's responsibility to check for a lock and even if a lock is detected, it can disregard it and perform IO anyway. The table below shows the locking semantics and mechanisms for NFSv3, NFSv4, and SMB.

Table 11 Locking mechanisms

Protocol	Advisory/mandatory	Mechanism
NFSv3	Advisory	Separate protocol (NLM)
NFSv4	Advisory or mandatory (default)	Embedded in the protocol
SMB	Mandatory	Embedded in the protocol

Due to the differences in the protocol specifications, the locking policy must be configured for the desired behavior on multiprotocol file systems. The protocol that is used and the locking policy setting determines whether or not a lock prevents IO:

- **Mandatory (default):** All IO must honor SMB and NFSv4 range locks. NFSv3 range locks never prevent IO since they are always advisory due to protocol specification.
- **Advisory:** NFSv3/v4/FTP IO bypasses all range locks. SMB bypasses NFSv4 range locks, but always honors SMB range locks due to protocol specification. Any lock requests continue to report lock conflicts.

The table below also shows which locks are honored for each protocol and locking policy setting in a chart format.

Table 12 Honored locks

Protocol	Mandatory (default)	Advisory
NFSv3	SMB + NFSv4	None
NFSv4	SMB + NFSv4	None
FTP	SMB + NFSv4	None
SMB	SMB + NFSv4	SMB

To configure the locking policy during file system creation, open the Advanced File System Settings menu as shown below.

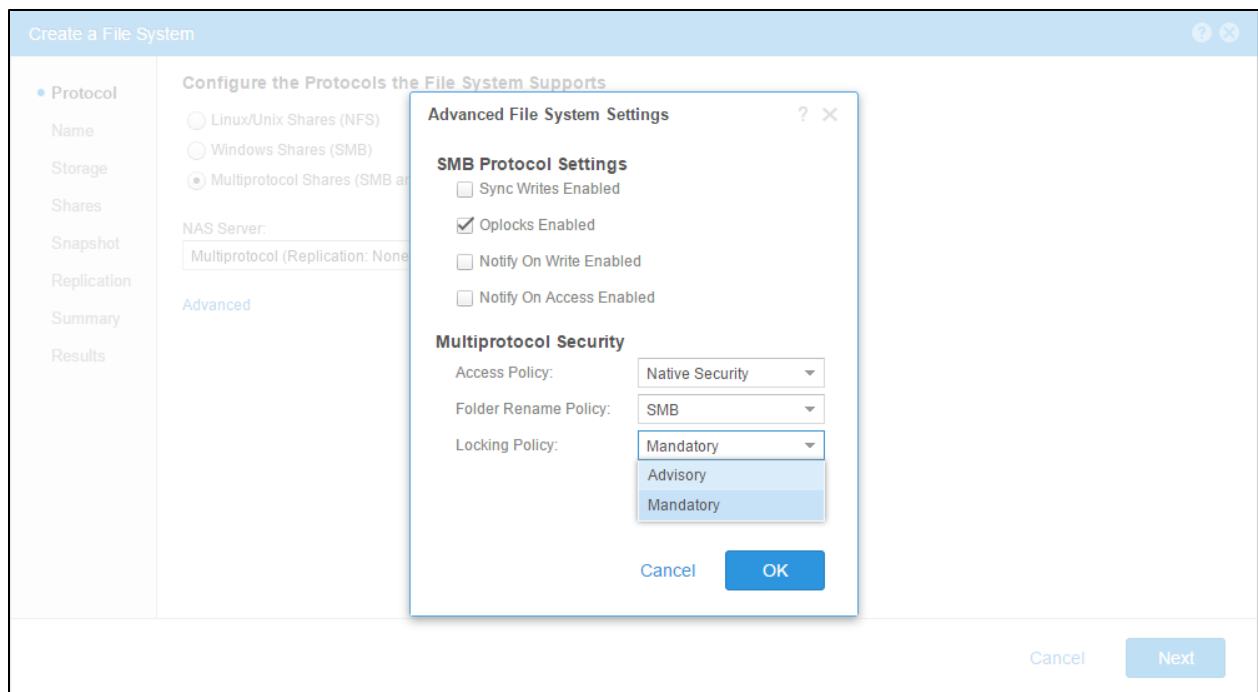


Figure 35 Locking policy

7.4.2 Folder rename policy

According to the SMB protocol specifications, renaming any directory that is located in the path of an open file is prohibited. For example, if C:\Folder1\Folder2\Folder3\File1.txt is opened by an SMB client, other clients are prevented from renaming any of the folders in the path leading up to File1.txt.

Clients using NFS or FTP do not have the same restriction. This is because SMB opens the entire path but NFS and FTP leverages file handles instead. Due to the differences between protocols, the Folder Rename Policy allows the storage administrator to configure the desired behavior on multiprotocol file systems. The

Folder rename policy settings are only invoked when attempting to rename a folder in a path of an open file. Renaming folders that do not have any open files in the path are always allowed.

The folder rename policy can be configured to:

- **Allowed:** All protocols can rename folders in the path of an open file without restrictions
- **SMB (default):** SMB protocol renames of a folder in the path of an open file are prohibited, but other protocols are allowed
- **Not Allowed:** No protocols are allowed to rename folders in the path of an open file

To configure the folder rename policy on an existing file system, navigate to the **File System Properties** → **Advanced** tab.

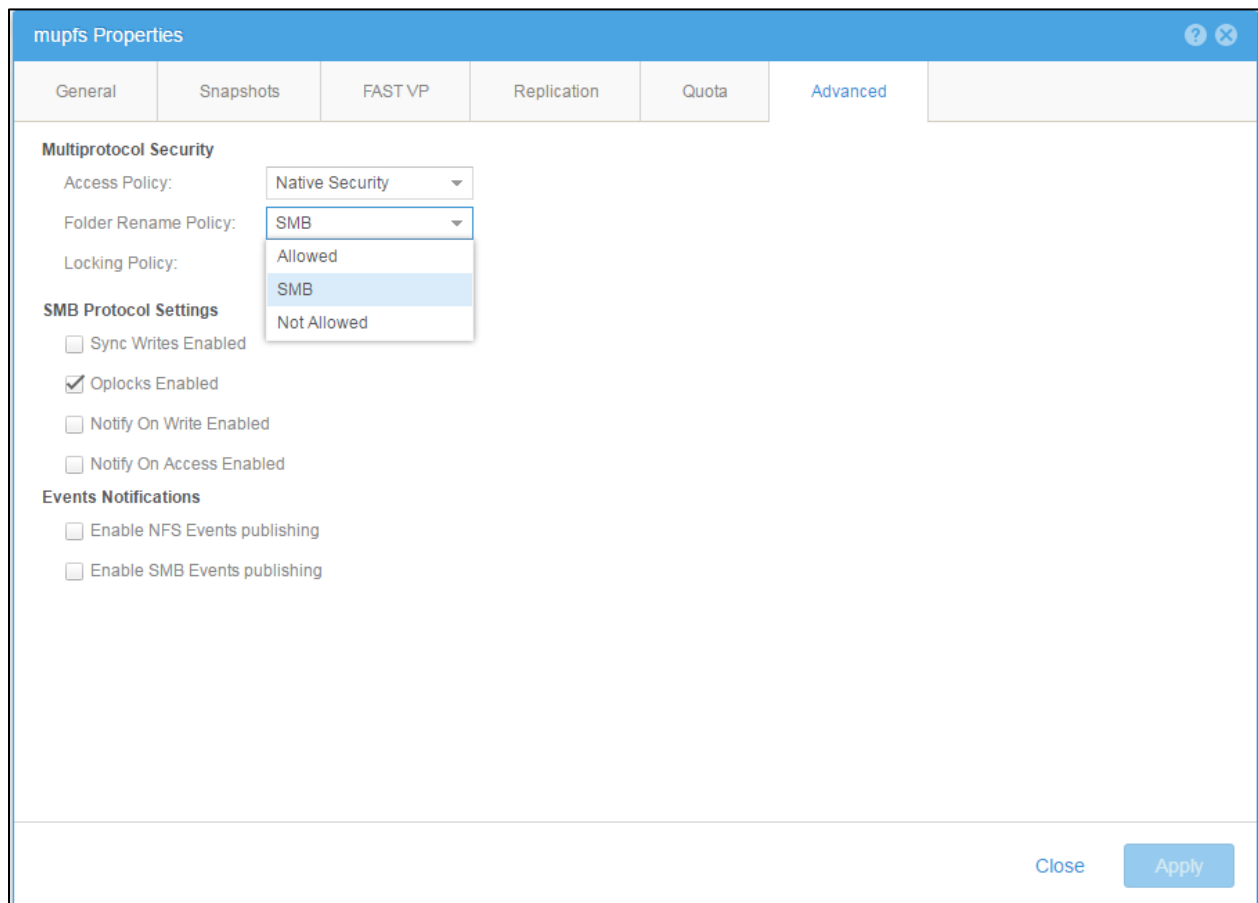


Figure 36 Folder rename policy

7.5 FTP and SFTP

Dell EMC Unity NAS servers and file systems also support access for FTP and/or SFTP (SSH File Transfer Protocol). SFTP is more secure since, unlike FTP, it does not transmit usernames and passwords in clear text. FTP and SFTP access is enabled or disabled individually at the NAS server level. Administrators can control the types of user accounts that can access files over FTP or SFTP, such as SMB, UNIX, and/or anonymous users. A home directory restriction option restricts access only to users who have existing home directories on the file system, however a default home directory can also be configured to allow all other users access to the file system when this restriction is applied. FTP and SFTP tracks and records connections and

file access for the NAS server. The audit logging settings also allow administrators to define the audit log file directory and the maximum size of audit log files.

For more granular control over access, FTP-enabled NAS servers support defining access control lists in the NAS server Properties page. Access can either be allowed or denied for a user-defined list of users, groups, and hosts in order to restrict FTP or SFTP access to only the desired users. However, users, groups, or hosts with restricted access to FTP or SFTP will still be able to access the NAS server and file systems over SMB or NFS as allowed by the ACLs or host access configurations for those protocols. The table below provides a list of FTP and SFTP protocol options.

Table 13 FTP/SFTP options

Protocol options	Default
Enable FTP	Disabled
Enable SFTP	Disabled
Allow SMB Users Access to the FTP/SFTP server	Enabled
Allow UNIX Users Access to the FTP/SFTP server	Enabled
Allow anonymous Users Access to the FTP server	Enabled
Home Directory Restriction	Disabled
Default Home Directory	/
Enable FTP/SFTP Auditing	Disabled
Directory of Audit Files	./etc/log
Maximum Size of Audit Files	512 KB

FTP access can be authenticated using the same methods as NFS or SMB. Once authentication is complete, access is then considered to be the same as SMB or NFS for security and permissions purposes. The method of authentication that is used depends on the format that is used for the username. If domain@user or domain\user is used, SMB authentication is used. For any other single username format, NFS authentication is used. SMB authentication uses the Windows Domain Controller while NFS authentication LDAP, NIS, or Local Files.

In order to use Local Files for FTP access, the **passwd** file must include an encrypted password for the users. This password is only used for FTP access. The Dell EMC Unity **passwd** file uses the same format and syntax as a standard UNIX system so that can be leveraged to generate the **passwd** file. On a UNIX system, use **useradd <user>** to add a new user and **passwd <user>** to set the password for that user. Then, copy the hashed password from the **/etc/shadow** file, add it to the second field in the **/etc/passwd** file, and upload it to the NAS server.

7.6 Internationalization

On Unity filesystems, the file and folder names are stored on disk using Unicode UTF-8 for encoding for the names, and Unicode names are presented for use on the network. NFSv4 clients only use UTF-8 in filenames. Modern SMB clients will use Unicode UTF-8 for file and folder names, but in SMB, part of the session negotiation can go to one of several Microsoft code pages to ensure compatibility over the network with file name encoding from older national Windows variants (which still use non-UTF-8 alphabets in their storage).

However, both NFSv3 and FTP protocols have no means of negotiating name-encoding in the protocol. Most modern Linux clients use UTF-8, and so are compatible with Unity. But some older Unix and Linux clients will use the file-name encoding that is specified in their system's 'locale' variables, and those alphabets can be very different from UTF-8. Most alphabets are compatible in the first 7-bits, or 127 characters, so you may not see errors on simple filenames which use only the old ASCII-character subset. Such clients often use one of the ISO-standard alphabets, usually one of the ISO-8859 series.

Clients using the 8859-x format can connect, read, and write filenames but some of their internationalized characters may not be displayed properly. In some cases it would be very difficult to have all the clients upgrade their 'locale' to a Unicode-using version. So in order to support these clients, a per-server translator must be configured to translate filenames from 8859-x (network format) to UTF-8 (disk format) since a per-session translation is not available. Dell EMC Unity OE version 4.3 includes the ability to configure a code page for NFSv3 and FTP clients. This is only required for NFSv3 and FTP clients since the protocol has no method of translation.

This setting does not affect NFSv4 and Windows clients as they always use UTF-8 or Unicode. It is highly recommended to configure this parameter prior to migrating or writing any NFSv3 data from another system on to Unity. Otherwise, the names of existing files containing internationalized characters are not recognized. If the parameter is changed, all NFSv3 and FTP clients connecting to this NAS server also need be configured for the same encoding type. Correctly configuring this setting also allows the other types of clients (NFSv4, SMB, or SFTP) to see the same filenames as the NFSv3 and FTP clients, including the correct internationalized characters.

The **vdm.codepage** parameter can be configured to:

- UTF-8 (default): Filenames are UTF-8 encoded
- 8859-1: Filenames are latin-1 encoded
- 8859-15: Filenames are latin-9 (latin-1 with the euro sign) encoded

For more information about NAS server parameters and how to configure them, reference the *Service Commands* document on [Dell EMC Support](#).

8 Features

Dell EMC Unity file systems and NAS servers support a wide range of features intended to optimize performance, improve efficiency and ensure important production data is protected in the event of a disaster. These features range from local snapshots, to remote replication and backup, to flash efficiency features, and more.

8.1 Data reduction

Data reduction helps reduce the total cost of ownership and increase the efficiency of the system. Dell EMC Unity OE version 4.2 includes compression support for file systems and VMware NFS datastores. On Dell EMC Unity OE version 4.3 and later, data reduction replaces compression and provides more space savings logic to the system with the addition of zero block detection and deduplication. In Dell EMC Unity OE version 4.5 and later, advanced deduplication is included as an optional feature to the data reduction algorithm. This provides the ability to reduce the amount of storage needed for user data by only keeping a single copy of a data block and removing all duplicates.

Data reduction and advanced deduplication reduce the amount of physical storage required to store a dataset, which can lead to cost savings. These savings are not only limited to the storage resource itself, but also on snapshots of those resources as well.

Data reduction and advanced deduplication occur inline between the system cache and the storage resource on an all flash pool. Data reduction and advanced deduplication that work at the protocol level is not supported. The requirements for data reduction on file resources are:

- Thin: The resource must be thinly provisioned.
- All flash pool: The resource must be provisioned from an all-flash pool. All-flash pools can be created on a Dell EMC Unity Hybrid Flash or All Flash System.
- Dell EMC Unity OE: File systems must be created on Dell EMC Unity OE version 4.2 or later and the system itself must currently be running Dell EMC Unity OE version 4.3 or later. Existing file systems created prior to Dell EMC Unity OE version 4.2 cannot have data reduction enabled. To achieve data reduction savings on older file systems, follow the procedure described in the *Dell EMC Unity: Compression for File* document on [Dell EMC Support](#).

In order to leverage advanced deduplication, the requirements for data reduction must first be met. In addition, the requirements for advanced deduplication must also be met:

- Dell EMC Unity OE:
 - Dell EMC Unity OE version 4.5 requires dynamic pools
 - Dell EMC Unity OE version 5.0 and later can use either dynamic or traditional pools
- Only the following models support advanced deduplication:
 - Dell EMC Unity 450F, 550F, and 650F
 - Dell EMC Unity XT 380, 380F, 480, 480F, 680, and 680F, 880, and 880F

For file resources that meet these requirements, data reduction and advanced deduplication can be enabled during creation and enabled or disabled at a later time. Data reduction and advanced deduplication can also be enabled on resources participating in replication sessions. The source and destination storage resources in a replication session are completely independent, and data reduction and advanced deduplication can be

enabled or disabled separately on the source and destination resource. The availability of enabling data reduction and advanced deduplication on a source and/or destination resource depends on the Dell EMC Unity OE version, system type, Pool configuration, and system model.

For more information about data reduction and advanced deduplication, reference the *Dell EMC Unity: Data Reduction* white paper on [Dell EMC Support](#).

8.2 Local protection

8.2.1 Snapshots

Dell EMC Unity snapshots are fully unified and use a common redirect-on-write technology between file and block. Dell EMC Unity snapshots are taken, treated, and scheduled the same way between block and file resources, resulting in a fully unified data protection experience. Unlike other file snapshot implementations, Dell EMC Unity file system snapshots do not require a separate volume to be set aside to accommodate snapped data. Instead, changes to snapped file system data are simply written to free space in the same storage pool. Snapshots can be read-only or read/write.

In the figure below, a snapshot is taken of a source file system containing data blocks A, B, C, and D. Afterward, when new data, D', is written to block D, the new data is redirected to a new location within the same pool and the data in block D is preserved as part of the snapshot. This works the same way when writing to snapshots that share data with the production file system. Unless the data is unique to the snapshot, attempting to overwrite a block will redirect the new block to a new location in the pool.




Figure 37 Unified snapshots

File snapshots may be restored through Unisphere on a file system level. You can also copy specific files from a snapshot back to the production file system by using the `.ckpt` directory in UNIX. However, note that some Operating Systems, such as Windows, may not be able to access the `.ckpt` due to directory caching. For Windows Operating Systems, the Previous Versions tab can be used instead to access the snapshots contents. The name of the `.ckpt` folder can also be customized by using the `cvfs.virtualDirName` parameter. For more information about NAS server parameters and how to configure them, reference the Service Commands document on [Dell EMC Support](#).

Dell EMC Unity also provides the ability to provide read/write access to file system snapshots to hosts and clients through shares. In order to do this, the administrator creates a new file system share using an existing snapshot. As a result, the snapshot data is exposed as a new share of the production file system, which may then be accessed as a normal share by file system hosts and clients.

For more information on Dell EMC Unity's unified snapshot technology, reference the *Dell EMC Unity: Snapshots and Thin Clones* white paper on [Dell EMC Support](#).

8.2.2 NDMP

Dell EMC Unity systems support 2-way and 3-way NDMP, allowing administrators to protect file systems by backing up to a tape library or other supported backup device. 3-way NDMP transfers the backup data over the network while 2-way NDMP transfers the data over Fibre Channel. 2-way NDMP eliminates backup data on the network by backing up directly to the backup device, potentially decreasing network congestion and reducing backup times. In order to use 2-way NDMP, the system must be running Dell EMC Unity OE version 4.4 or later. A 2-way NDMP configuration is shown in the figure below.

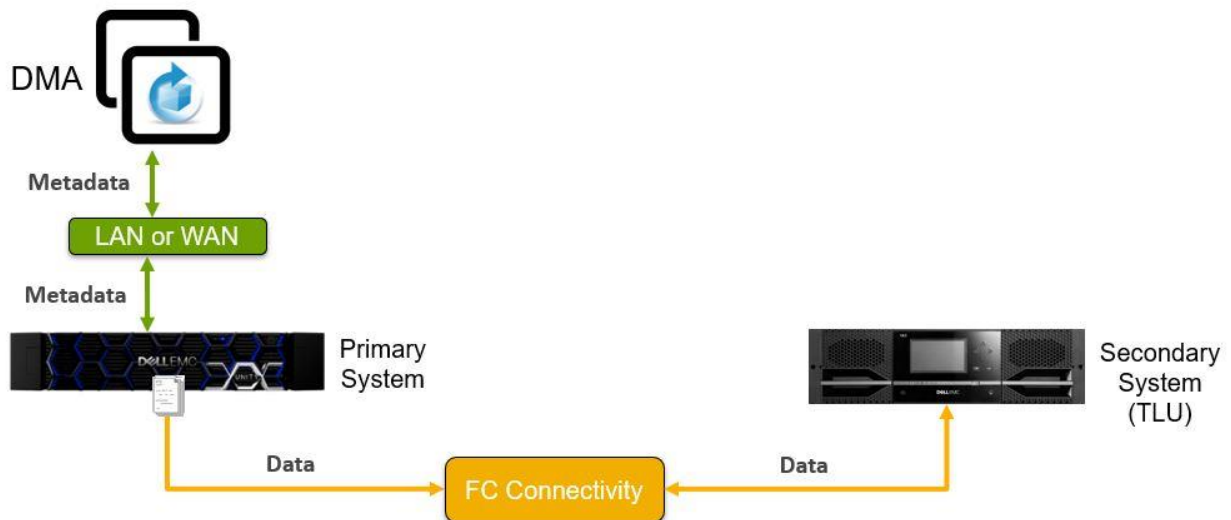


Figure 38 2-Way NDMP

When configuring 2-way NDMP, the backup device should be connected to a switch and zoned to the Fibre Channel ports on the Dell EMC Unity system. Directly connecting the backup device to the storage system is not supported. When cabling and zoning the system, using the Synchronous Replication port, which is the first Fibre Channel port on the system, for the backup device is not supported.

Dell EMC Unity supports taking NDMP full backups, incremental backups, restores, and tape cloning. Both dump and tar backups are supported but volume based backups (VBB) are not supported. The backup application can specify the parameters below when running an NDMP backup. It is recommended to enable all of these parameters when running an NDMP backup.

- HIST: Allows the backup application to request the file history from the storage system
- UPDATE: Allows the backup application to request the file history for incremental backups
- DIRECT: Enables the ability to restore a single file from a backup
- SNAPSURE: Allows the backup application to request a snapshot of the file system for backup purposes

Combining NDMP, local snapshots, and remote protection enable Dell EMC Unity storage systems to be deployed with a wide array of data protection capabilities, including the ability to replicate to or from multiple arrays in a multisite topology. In addition, the NDMP backups can be taken on the destination NAS server, alleviating the backup load from the production system.

8.3 Remote protection

Remote protection provides business continuity by copying the data to a remote system, enabling data access in situations where the primary system is no longer accessible. Dell EMC Unity supports both synchronous (MetroSync) and asynchronous replication for file and block resources.

In situations where one system is running a newer version of Dell EMC Unity OE than the other, replication is supported. However, if there is a new feature that is in use but is not available on the peer system, the replication session may not be supported and may fail to configure. For existing replication sessions, you may be prevented from enabling some of these new features. Some examples of this include CEE CEPA, IP multi-tenancy, and the folder rename and locking policies. In these situations, upgrade the older system to be on the same Dell EMC Unity OE as the newer system before configuring replication.

8.3.1 MetroSync

MetroSync is available on systems running Dell EMC Unity OE version 4.4 or later. This feature provides the ability to create remote synchronous replication sessions for file storage resources including NAS servers, file systems, and VMware NFS datastores. Synchronous replication is a zero recovery point objective (RPO) data protection solution which ensures each block of data stored locally is also saved to a remote image before the write is acknowledged back to the host. This ensures that in the event of a disaster, there is zero data loss. In synchronous replication solutions, there are trade-offs. As each write needs to be saved locally and remotely, added response time occurs during each transaction. This response time increases as distance increases between remote images. Synchronous replication has a distance limitation based on latency between systems. This limitation is generally 60 miles or 100 kilometers between sites. To support synchronous replication, the latency of the link must be less than 10 milliseconds.

Synchronous replication uses the first Fibre Channel (FC) port on the system to replicate both the NAS server and file system data. The synchronous replication management virtual port is used to send management and orchestration commands between systems. Since there is no Fibre Channel support on Dell EMC UnityVSA systems, synchronous replication cannot be configured on the virtual storage appliance.

Synchronous replication requires two separate physical Unity systems, meaning it cannot be used to replicate file resources locally within the same system. Both the source and the destination systems must be running Dell EMC Unity OE version 4.4 or later in order to support synchronous replication.

In order to synchronously replicate a file resource, the associated NAS server must be synchronously replicated first. After this is configured, synchronous replication can be configured on its associated file systems. When MetroSync is configured, the following functionality is also available:

- Snapshot replication: Synchronous replication of read-only snapshots to the destination system. Snapshot replication occurs automatically while the session is in sync.
- Snapshot schedule replication: Synchronous replication of snapshot schedules to the destination system. This ensures the destination system has the same snapshot schedules applied as the source system in case of failover.
- Cabinet level failover: Single command to initiate a simultaneous failover of all synchronously replicated NAS servers and their associated file systems in parallel. This should only be used if the source system is offline and unavailable.
- Asynchronous replication to a third site: With synchronous replication configured between the two primary sites, this enables the ability to add asynchronous replication to a third site for backup purposes. In case of failover between the two primary sites, the asynchronous replication sessions to the third site can be incrementally restarted without requiring a full resync.

In Dell EMC Unity OE version 4.5, MetroSync Manager is available. This is an optional Windows application that monitors the systems participating in synchronous replication for critical failures, such as a power outage. If a failure is detected, MetroSync Manager initiates a cabinet level failover to simultaneously failover all synchronously replicated NAS servers and their file systems to the peer system in parallel. Without MetroSync Manager, manual intervention is required to initiate the failover process which could lead to periods of data unavailability.

For more information on Dell EMC Unity MetroSync and MetroSync Manager, reference the *Dell EMC Unity: MetroSync* white paper on [Dell EMC Support](#).

8.3.2 Asynchronous replication

Asynchronous replication is primarily used to replicate data over long distances, but also can be utilized to replicate file resources between pools within the same system. Asynchronous replication does not impact host I/O latency as host writes are acknowledged once they are saved to the local storage resource. Because write operations are not immediately replicated to a destination resource, all writes are tracked on the source. This data is replicated during the next synchronization.

Asynchronous replication introduces the concept of a recovery point objective (RPO). RPO is the acceptable amount of data, measured in units of time, which may be lost due to a failure. This delta of time also affects the amount of data which needs to be replicated during the next synchronization, and the amount of potential data loss if a disaster scenario were to occur.

By leveraging the unified snapshots technology to preserve point-in-time file and block data, Dell EMC Unity is able to provide unified local and remote asynchronous replication using the same technology for file and block resources. Dell EMC Unity supports asynchronous replication down to a 5-minute RPO for NAS servers and their file systems.

Asynchronous Replication is supported on all Dell EMC Unity systems for both file and block resources. The supported systems include the All Flash models, the Hybrid models, and the Dell EMC UnityVSA.

With the release of Dell EMC Unity OE version 5.0, asynchronous replication can be configured in advanced topologies. This allows for configurations such as fan-out and cascading replication at the file resource level along with the corresponding NAS server. These configurations give the ability to replicate and store the same dataset on multiple systems provides additional data protection and enables use cases such as content distribution. This feature does not support synchronous replication and is not available for block resources.

In order to use this feature, all systems in the topology must be running Dell EMC Unity OE version 5.0 or later. Although, replication between Dell EMC Unity OE 5.0 and 4.x is still supported in a one-directional configuration. Aside from that, this feature is available on both physical and virtual Dell EMC Unity systems. The following advanced topologies are supported:

File resource level:

- One-directional
 - A single file resource replicating to a single destination resource
 - Example: A → B
- Fan-out
 - A single file resource replicating to up to four different destination systems

- Example: $A \rightarrow B$ and $A \rightarrow C$

- Cascade
 - A single file resource replicating to a second system and from there, replicating to a third system
 - Example: A → B → C
- Mixed
 - Leveraging a combination of both cascade and fan-out, or fan-out and cascade
 - Example: A → B and B → C and B → D

In the figure below, a topology has been constructed that showcases both cascade and fan-out replication configurations.

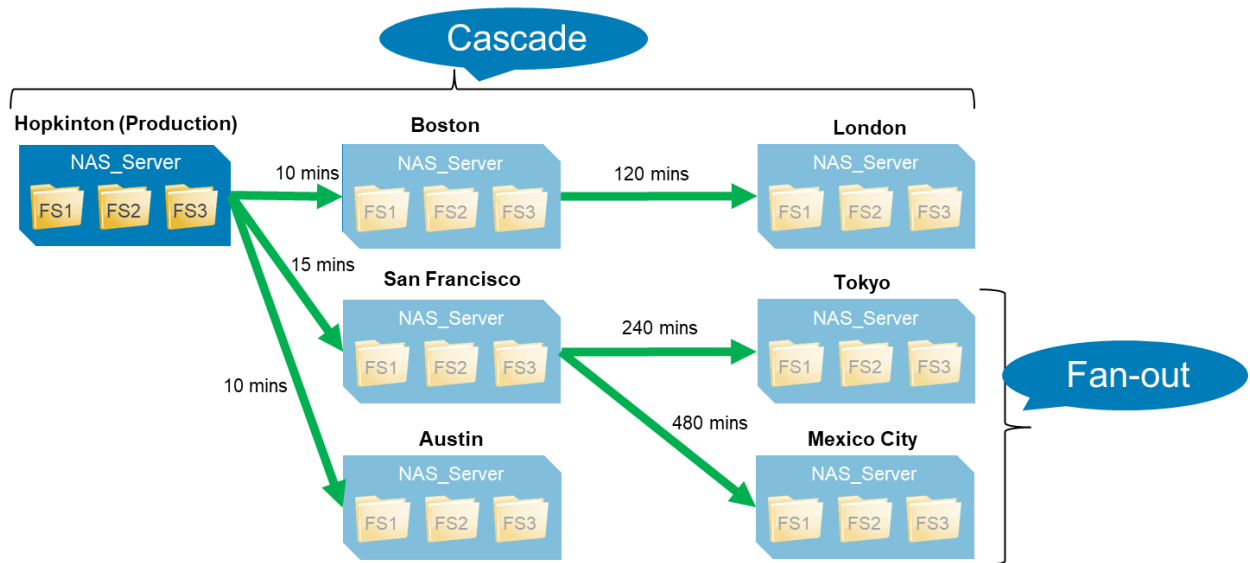


Figure 39 Advanced replication topologies

Refer to the “Advanced file remote replication” section of the *Dell EMC Unity: Replication Technologies* white paper for more details on the capabilities of this feature.

Prior to Dell EMC Unity OE version 5.0, there can only be a single source system and single destination system for a given file resource. However, there could be multiple resources on a single system replicating to different target systems as shown in the figure below.




Figure 40 Replication

Dell EMC Unity provides the ability to perform operations such as failover, failback, pause, and resume on individual NAS servers and file systems. For example, in order to initiate a failover, you must first fail over the NAS server and then fail over the individual file systems afterwards to enable access on the destination system. Dell EMC Unity OE version 4.2 includes an enhancement that enables group operations. The group operation automatically fails over all of the associated file systems if a failover is initiated on the NAS server. Group operations can be used for failover, failover with sync, failback, pause, and resume. These operations remain available at the individual file system level, but any operation applied at the NAS server level is a group operation. All other replication related operations such as create, sync, delete, and modify remain available only as individual operations. The figure below shows the failover group operation.

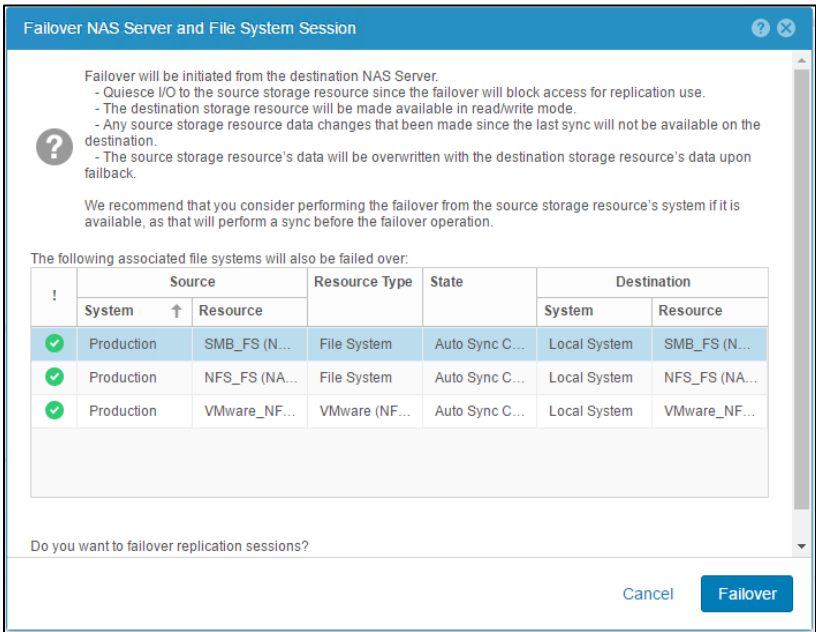


Figure 41 Failover group operation

The NAS server and file systems' replication sessions must be in a healthy state for group operations to properly function. If a NAS server replication session is in a paused or error state, all group operations are prohibited. If any of the associated file system replication sessions are in a paused, error, or non-replicated state, the group operation skips those particular sessions and an alert is generated. Also, do not perform a group operation at both sides of a replication session at the same time. This action is not prohibited by the storage system, however, a group operation performed at the same time at both sides of a replication session can cause the group replication session to enter an unhealthy state.

In order to minimize downtime during a failover with sync group operation, it is recommended to perform a manual sync on the NAS server and all associated file systems before starting the group operation. The manual sync operation performs an incremental update so that the majority of changes are replicated to the destination. Wait until all of the sync operations are complete prior to starting the failover with sync operation. Because of the previous incremental sync, any remaining deltas are small, which enables the failover with sync group operation to finish quickly and for downtime to be minimized.

Dell EMC Unity OE version 4.2 also includes the ability to replicate file and block snapshots along with the primary storage resource. Some benefits include the ability to configure a different retention policy on the destination for compliance purposes, improving cost efficiency by freeing up capacity from snapshots on the source system, and improvising resiliency by storing snapshots in a different fault domain. This feature requires both the source and destination system to be running Dell EMC Unity OE version 4.2 or later. Snapshot replication from Dell EMC Unity OE version 4.4 or earlier to Dell EMC Unity OE version 5.0 is not supported, it is recommended for users to first upgrade to Dell EMC Unity OE version 4.5 to replicate to a version 5.0 system. Also, an asynchronous replication session must be created on the primary storage resource. This feature can be enabled during creation of the replication session or at any time afterwards.

This feature works with both manually initiated and snapshots taken by the integrated snapshot scheduler. A replicated snapshot retains the same properties and attributes as the source snapshot. This includes the name, description, creation time, taken by, and so on. For file snapshots, only read-only snapshots are eligible for replication. Read/write snapshots that are used for shares and exports cannot be replicated.

For more information on Dell EMC Unity native asynchronous replication, reference the *Dell EMC Unity: Replication Technologies* white paper on [Dell EMC Support](#).

8.3.3 RO proxy NAS servers

Dell EMC Unity OE version 4.3 includes support for read-only proxy NAS servers. This provides the ability to access all the file system and snapshot data on the destination NAS server through SMB and NFS. There is no ability to write to the file systems or snapshots using the proxy NAS server, even if the snapshot is read/write.

The figure below shows the proxy NAS server configuration.

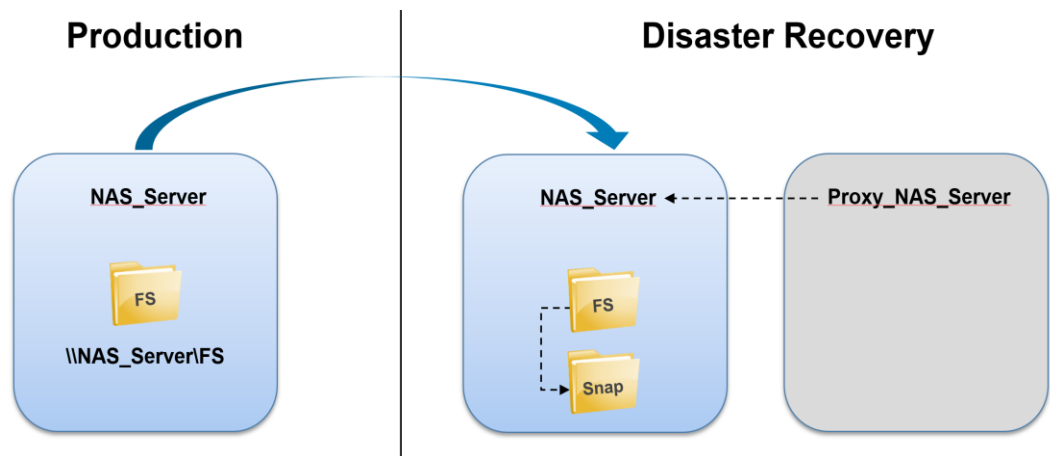


Figure 42 Proxy NAS servers

This feature works with both asynchronous and synchronous replication. Although it may be possible to directly access the file system data using the proxy NAS server, it is recommended to use this feature to access data residing on snapshots. This is due to the fact that the destination file system is still being actively replicated. For asynchronous replication, there may be instances where the destination file system needs to be frozen due to a replication sync. For synchronous replication, the destination file system is not fully mounted and is not accessible.

Each proxy NAS server can be configured to provide access to one or more NAS servers' data. All the NAS servers' file systems and their snapshots are displayed when connecting to the proxy NAS server as shown in the figure below. Due to this, the user must be part of the local administrators group for SMB or root for NFS to access the data.

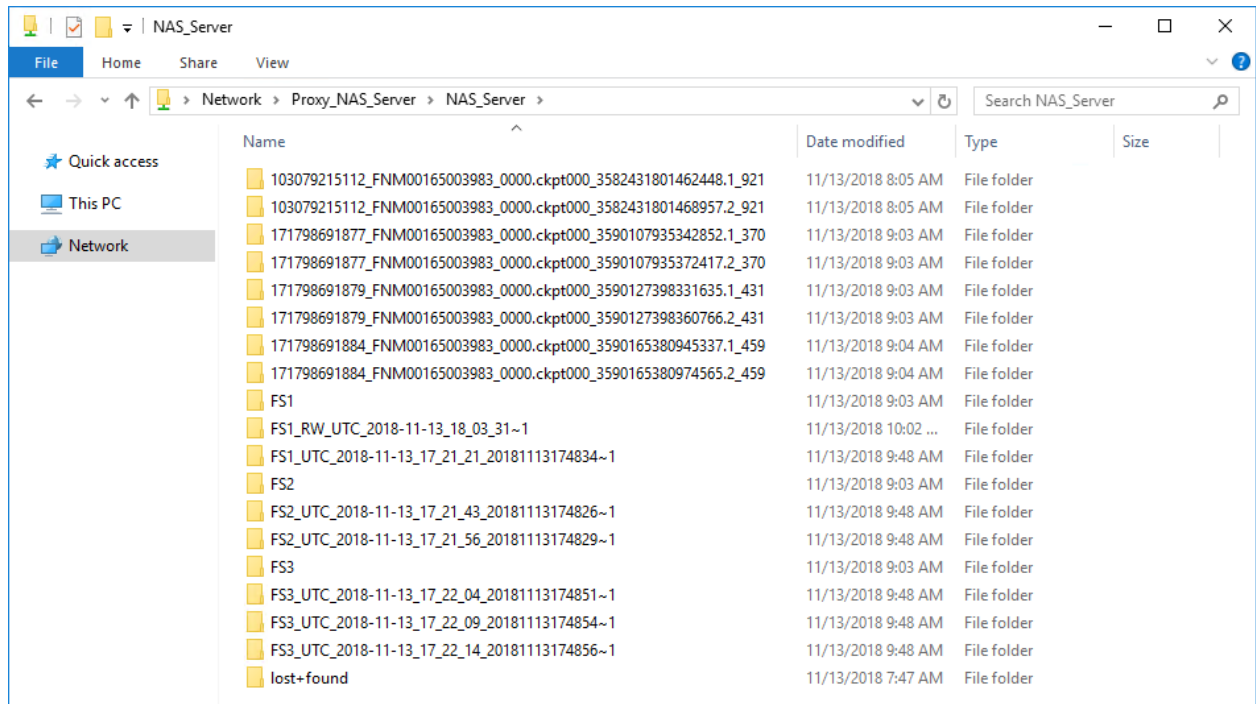


Figure 43 Proxy NAS server

In order to create a proxy NAS server, create a new NAS server on the system with an interface, the appropriate protocols, and configure the appropriate services such as DNS and LDAP. The new proxy NAS server should be configured the same way as the NAS servers that it is providing access to such as protocols, tenants, and so on. The new proxy NAS server must reside on the same SP as the NAS server that it will be providing access to.

In order to designate the new NAS server as a proxy NAS server, a CLI service command must be used. SSH in to the system and run the **svc_nas <Proxy_NAS_Server> -proxy -add <Target_NAS_Server> command** where **<Proxy_NAS_Server>** is the name of the new proxy NAS server you just created and **<Target_NAS_Server>** is the name of the destination NAS server you want users to access. For NFS access, also include **-NFSRoot <Allowed_Nodes>** option to specify the nodes that should have access over NFS. To view the proxy NAS server configuration on the system, run the **svc_nas ALL -proxy -show** command.

For example, run **svc_nas Proxy_NAS_Server -proxy -add NAS_Server -NFSRoot ip=10.10.10.10** to configure the proxy NAS server for NFS access and limit access to client IP 10.10.10.10. Run **mount Proxy_NAS_Server:/NAS_Server /mnt** on the host that is provided access or **\\Proxy_NAS_Server\NAS_Server** from a SMB client to mount the proxy NAS server and view the contents.

For more information on configuring Proxy NAS servers, reference the *Dell EMC Unity: DR Access and Testing* document on [Dell EMC Support](#).

8.3.4 RW SMB proxy shares

Dell EMC Unity OE version 4.5 introduces the ability to create SMB shares for writeable and read-only snapshots on the destination NAS server. This feature is designed to enable DR testing without any impact to the ongoing replication session. It allows customers to confirm that an application can be brought online and write to a share hosted on the destination system. This feature works with both asynchronous and synchronous replication. This feature leverages a proxy NAS server and proxy share created on the destination system to provide access to the snapshot, as shown in the figure below.




Figure 44 SMB RW proxy shares

In contrast to the read-only proxy NAS server feature, this feature allows any domain user to access the share and is not limited to Administrators or root. This is because each share points to a specific snapshot, as opposed to the entire contents of the NAS server. The proxy share can be configured to point to either a RO or RW snapshot that exists on the destination file system. If a RW snapshot is selected, then the client can write to the share.

To configure a proxy share, a new proxy NAS server must be created on the destination Dell EMC Unity system. The NAS server must reside on the same SP it is providing access to and must be joined to the same SMB domain as the destination NAS server. If these requirements are met, a single proxy NAS server can be used to access data on one or more destination NAS servers.

Once the proxy NAS server is configured, SMB shares can be created for snapshots. These special proxy SMB shares can only be configured and managed by using the **svc_nas** command. Once created, these shares are not visible through normal interfaces such as Unisphere, UEMCLI, or REST API. These shares also do not count towards the system limits and there is no hard limit on how many proxy SMB shares can be created.

To create a proxy SMB share, use the **svc_nas <Proxy_NAS_Server> -proxy_share -add <Target_NAS_Server> -share <Share_Name> -path <Snapshot_Path>** command, where:

- **<Proxy_NAS_Server>**: Name of the proxy NAS server
- **<Target_NAS_Server>**: Name of the NAS server it is providing access to
- **<Share_Name>**: Name of the share that the client uses to mount
- **<Snapshot_Path>**: Path to the RO or RW snapshot, usually this is the name of the snapshot prefixed with a /

For example, **svc_nas Proxy_NAS_Server -proxy_share -add Prod_NAS_Server -share FS -path /UTC_2018-09-13_15:31:25**. Once this is created, any domain user can access the snapshot by mapping the UNC path **\\Proxy_NAS_Server\FS**, as shown in the figure below.

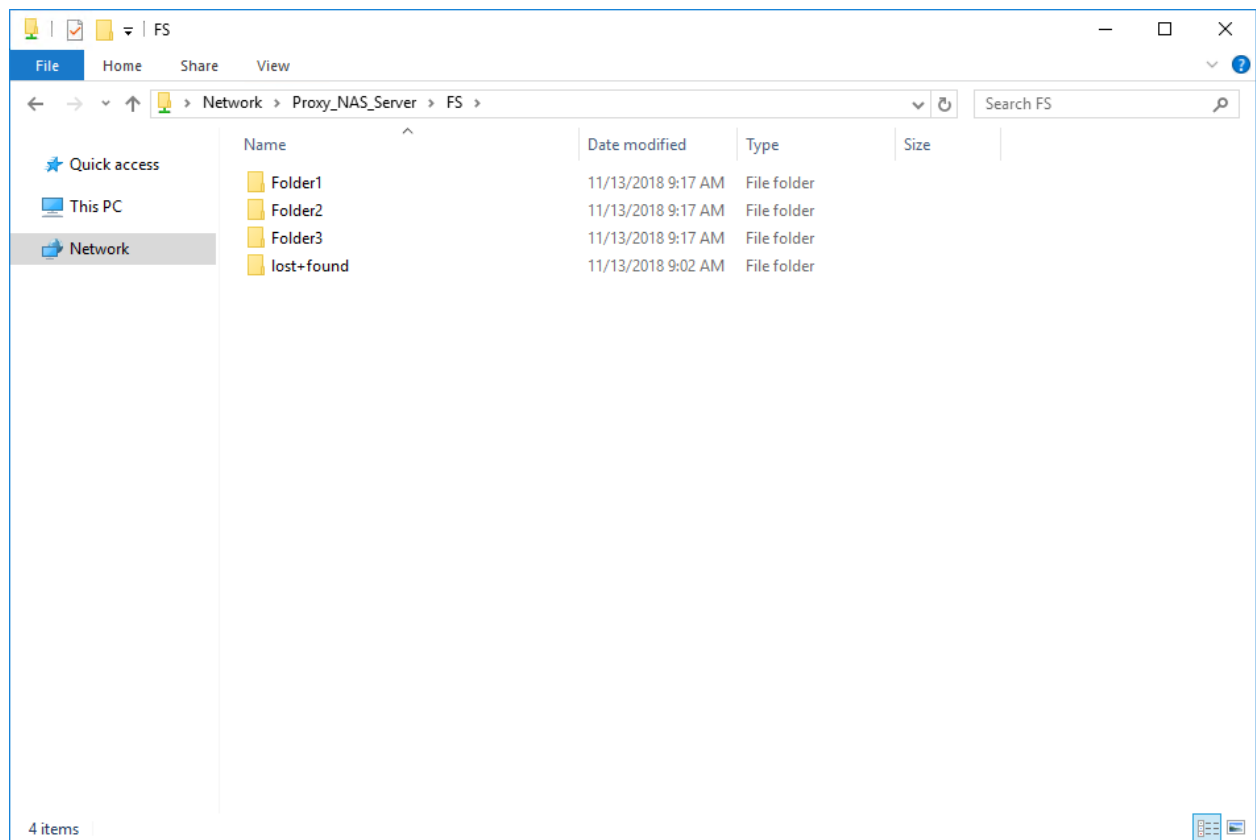


Figure 45 SMB proxy NAS share

For more information on configuring SMB proxy NAS shares, reference *the Dell EMC Unity: DR Access and Testing* document on [Dell EMC Support](#).

8.3.5 Interfaces

When using replication, the interfaces and services on the destination system can be configured with overridden IP addresses. This enables the use of separate IP addresses if there is a failover to the DR site. These can be pre-configured on the DR NAS server, so production can resume as soon as the failover process is complete. This is important for customers that have different network schemas between their production and DR sites. In addition, backup and test interfaces can also be configured on the NAS server to enable NFS access for backup and DR testing purposes. These interfaces can be active on either the source or the destination NAS server, but they are not replicated in the replication session.

When replicating from one system to another, it is important to ensure the port configuration matches on both systems. If link aggregation and/or FSN are used, ensure the link aggregation and/or FSN also exists on the destination system. If the specified port, link aggregation, or FSN is not found on the destination system, the interfaces on the destination NAS server are created without a port assignment. You must then override the IP addresses on the destination NAS server to assign them to a valid port. Otherwise, data access becomes unavailable in the event of a failover.

8.4 FAST technology

Dell EMC Unity also provides performance optimization through FAST™ Cache and FAST VP on Dell EMC Unity hybrid systems. Since file systems are provisioned out of the same Unified pools, FAST VP can be applied directly to individual file systems themselves, rather than second hand through file LUNs forming a file storage pool, resulting in greater granularity and efficiency. Because of this, FAST can be leveraged more effectively and granularly on Dell EMC Unity file systems. One such example is FAST VP tiering policies, which can be assigned on a per file system basis, allowing each individual file system to be assigned one of the following policies: highest available tier, start high then auto-tier (default), auto-tier, or lowest available tier. As a consequence it is possible to have a single unified storage pool consisting of many file systems with different tiering policies, allowing more flexibility when configuring a system to take advantage of FAST VP's powerful tiering capabilities. Furthermore, this tiering will be equally as effective for both file or block resources, because the resources now share a unified storage pool. Similarly, FAST Cache benefits LUNs and file systems equally, ensuring that appropriate transactional file workloads see a large performance improvement from even a small amount of flash capacity. For more information on FAST VP and FAST Cache, reference the *Dell EMC Unity: FAST Technology Overview* white paper on [Dell EMC Support](#).

8.5 Custom file alert thresholds

Dell EMC Unity file systems and NFS datastores support file alert thresholds, which generate warning and error level alerts as space utilization increases. When a file system or NFS datastore passes 75% utilization, a warning level alert is generated. The utilization percent is calculated based on used space, and savings achieved by data reduction and advanced deduplication are not factored in. If the utilization of the resource continues to increase, an error level alert is generated at 95% utilization. This ensures that administrators are able to easily identify and address any capacity related issues with their file systems or NFS datastores before they run out of space to save data.

With the release of Dell EMC Unity OE version 5.0, custom file alert thresholds are introduced. This feature allows administrators to configure the utilization percentages that a file system or NFS datastore will generate an informational (disabled by default), warning, and/or error level alerts. Thresholds can be configured

between 0-99%, and the utilization percent must increase with each subsequent alert severity level. A setting of 0% disables the alert. The default settings are displayed below:

- Information: 0% (disabled)
- Warning: 75%
- Error: 95%

Custom file alert thresholds can be modified by navigating to the Properties page of a file system or NFS datastore and selecting the **Capacity Alarm Setting** link.

8.6 Common Event Enabler

Common Event Enabler (CEE) is a software package that runs on a Windows or Linux server. CEE consists of two parts:

- **Common Antivirus Agent (CAVA)**: Allows antivirus engines to scan files stored on file systems
- **Common Event Publishing Agent (CEPA)**: Allows applications to receive file event notifications

8.6.1 CAVA

CEE CAVA provides an antivirus solution to SMB clients by using third-party antivirus software to identify and eliminate known viruses before they infect files on the storage system. Windows clients require this to reduce the chance of storing infected files on the file system and protects them if they happen to open an infected file. This antivirus solution consists of a combination of the Dell EMC Unity system, CEE CAVA agent, and a third-party antivirus engine. CEE CAVA is enabled on a per NAS server basis in the **NAS server Properties → Security → Antivirus** page.

The Dell EMC Unity system monitors events and triggers the AV engine to initiate a scan when necessary. Some of the possible event triggers include file renames, modifications, and first reads. While a file is being scanned, access to the file from any SMB client is temporarily blocked. The CEE CAVA solution is for clients running the SMB protocol only. If clients use the NFS or FTP protocols to create, modify, or move files, the CEE CAVA solution does not scan these files for viruses.

CEE CAVA can be customized depending on your specific needs. It has the ability to scan specific file extensions, exclude specific file extensions, configure the maximum file size to be scanned, configure the desired behavior if the AV server goes offline and more. To ensure that file scanning is maintained if an AV server goes offline or cannot be reached, you should configure at least two CAVA servers.

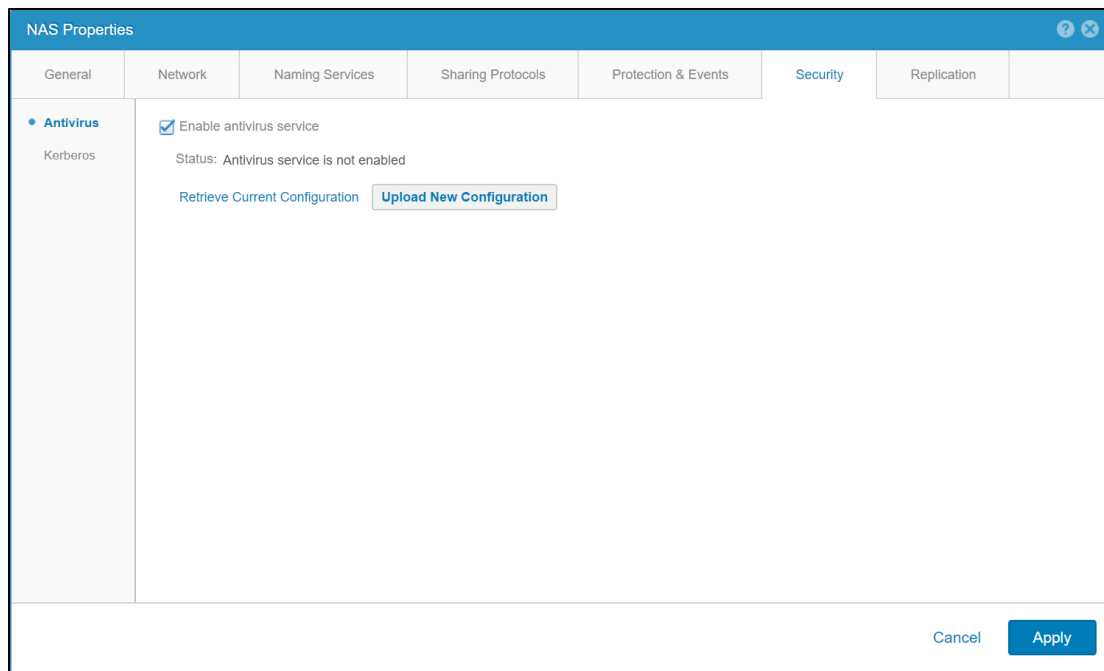


Figure 46 CEE CAVA

For a list of supported AV engines or for more information on how to configure CEE CAVA, reference *the Dell EMC Unity Simple Support Matrix* and *the Dell EMC Unity Family: Configuring Hosts to Access SMB File Systems* document on [Dell EMC Support](#).

8.6.2 CEPA

The CEE CEPA ecosystem consists of dozens of applications that are designed to process SMB and NFS file and directory event notifications. Starting with Dell EMC Unity OE version 4.1, CEE CEPA also works on Dell EMC Unity systems. Some of the common uses for CEE CEPA include file auditing, centralized quota management, search/indexing, and as an enabler for other use cases such as RabbitMQ. An example of file auditing is logging an event any time a file or directory is created, renamed, or deleted. To configure CEE CEPA, CEE and a third-party file event software must be installed and configured in the environment.

There are three types of events that can be logged:

- **Pre Events:** Events are sent to the CEPA server for approval prior to being executed
- **Post Events:** Events are sent to the CEPA server after they occur for logging or auditing purposes
- **Post Error Events:** Error events are sent to the CEPA server after they occur for logging or auditing purposes

CEE CEPA must be enabled and an Event Publishing Pool must be created on the NAS server, which defines the CEPA servers and the specific events that trigger notifications. This can be found in the **NAS server Properties → Protection & Events → Events Publishing** page. You can create up to three CEPA Pools per NAS server, but only one of these can be configured for Pre Events.

CEE CEPA also has policy settings that can be configured to determine the desired behavior if all the CEPA servers are unavailable:

- **Pre Events Failure Policy:** The behavior if all CEPA servers go down that are configured for pre events
 - **Ignore (default):** Assume all events are approved
 - **Deny:** Deny events that require approval until the CEPA server comes back online
- **Post Events Failure Policy:** The behavior if all CEPA servers go down that are configured for post events
 - **Ignore (default):** Continue operating, losing any events that occurred while the CEPA server is down
 - **Accumulate:** Continue operating and save events to a local circular buffer (500MB)
 - **Guarantee:** Continue operating and save events to a local circular buffer (500MB) and once it's full, deny access
 - **Deny:** Deny access to file systems until the CEPA server comes back online

Once the NAS server is configured, enable events publishing on the file systems that you want to receive events from. This can be done under the **File System Properties → Advanced** tab. Depending on the type of file system, you can enable NFS and/or SMB events publishing.

When a host generates an event on the file system over SMB or NFS, this information is forwarded to the CEPA server over an HTTP connection. The CEE CEPA software on the server receives and publishes this event, enabling it to be processed by the third-party software. This enables the CEPA server to provide many additional services with information being sent from multiple systems.

For more information on CEE CEPA, reference the *Using the Common Event Enabler* document on [Dell EMC Support](#).

8.7 Cloud Tiering Appliance

Cloud Tiering Appliance (CTA) enables the ability to tier data from Dell EMC Unity to another location based on user-configured policies. An example of this includes moving any files that are larger than 50 MB and has not been accessed in 30 days to the cloud. After a file is moved off the primary storage, an 8 KB stub is left which points to the actual location of the data. Any requests to read the data that has been moved can be configured as a pass-through, partial recall, or full recall from the cloud. From the end user point of view, this process is seamless since the stub resembles the actual file and the data continues to be accessible on demand. Archived files are marked with the offline file icon in Windows Explorer, so they can be easily identified.

The CTA interface is shown in the figure below.

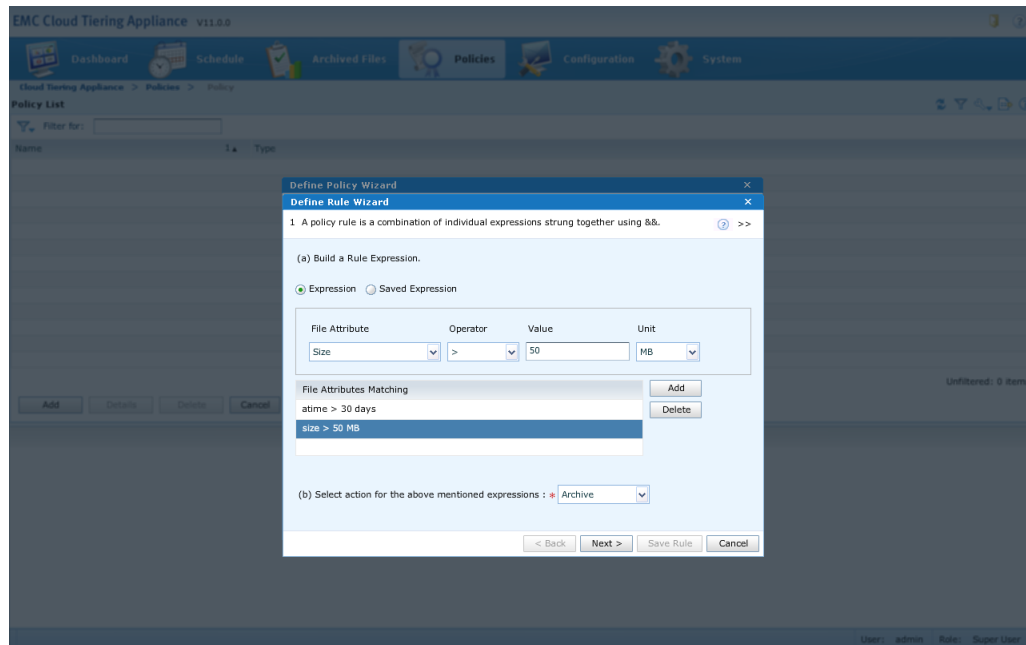


Figure 47 Cloud Tying Appliance

With Dell EMC Unity OE version 4.2 and CTA version 12 SP1, the CTA file migration feature is supported. This feature can be used to migrate file data from NetApp or VNX to Dell EMC Unity. The destination file system must have multiprotocol-enabled and needs enough capacity to hold the entire dataset from the source. For example, if deduplication and compression is enabled on the source file system, the destination file system is rehydrated so that must be considered.

Starting with Dell EMC Unity OE version 4.2 and CTA version 12, support for tiering of block snapshots to the cloud is available. This provides the ability to archive and restore block snapshots to and from the cloud. This works with both LUNs and consistency groups and leverages iSCSI connectivity between CTA and the Dell EMC Unity system. Similar to file tiering, tiering of block snapshots is also based on policies and schedules created by the administrator. Once a snapshot is archived to the cloud, it can be deleted off the Dell EMC Unity system to free up storage. Snapshots can also be restored from the cloud on to a new storage resource on the Dell EMC Unity system.

Starting with Dell EMC Unity OE version 4.1 and CTA version 11, CTA support for Dell EMC Unity is available. When using CTA with Dell EMC Unity as the source, only tiering to Virtustream, Azure, and S3 cloud repositories are supported. When tiering to a cloud repository, CTA can also leverage compression and/or encryption. Although CTA has the ability to tier other Dell EMC storage, this is not supported when Dell EMC Unity is configured as the source.

There are two deployment options: CTA and CTA-VE (Virtual Edition). CTA comes as an ISO image that can be installed on a physical server. There is no option to purchase a physical CTA appliance from Dell EMC so a customer-supplied server must be used. CTA/VE is installed on ESXi and runs as a VM. Both options include CTA-HA (High Availability) to provide redundancy for recall operations. Recall requests can also be configured to use both CTA and CTA-HA in a round-robin fashion to load balancing.

The minimum requirements for each are:

- CTA
 - 4-Core CPU
 - 16 GB RAM
 - 1 TB Capacity
 - 2 x 1Gb or 10Gb Interfaces
- CTA-VE
 - 4 vCPUs
 - 16 GB RAM
 - 1 TB Capacity
 - 2 x 1Gb or 10Gb Virtual Interfaces
- CTA-HA
 - 4 vCPUs
 - 4 GB RAM
 - 100 GB Capacity
 - 2 x 1Gb or 10Gb Virtual Interfaces

Once CTA is deployed and configured, navigate to **Configuration → Common API Settings**. Provide the Dell EMC Unity management credentials and set a DHSM password. When a Dell EMC Unity system is added as a source file server, this information is used to configure DHSM automatically on the NAS server. Alternatively, you can still configure it manually by enabling DHSM and setting a password under **NAS server Properties → Protection & Events → DHSM**. If multiple Dell EMC Unity systems are configured on a single CTA, the Dell EMC Unity management credentials and DHSM password has to be same across all of the Dell EMC Unity systems.

In CTA, add Dell EMC Unity as a source file server by providing details such as NetBIOS name, Dell EMC Unity management IP, authentication details, and the DNS name of CTA. To configure a cloud repository, enable the Cloud Callback Daemon and provide information such as the endpoint, access key, secret key, and bucket. Verify connectivity for both Dell EMC Unity and the cloud repository to ensure the configuration is correct. Then, create policies to tier data depending on specific attributes and schedules to run these jobs automatically.

For more information on CTA, reference *the Dell EMC Unity: Cloud Tiering Appliance (CTA) white paper* and *Cloud Tiering Appliance: Getting Started Guide* on [Dell EMC Support](#).

8.8 File import

Dell EMC Unity OE version 4.1 includes Native File Import for NFSv3 and Dell EMC Unity OE version 4.2 includes Native File Import for SMB. This feature provides a native option for file migration to Dell EMC Unity from a VNX1 or VNX2 system.

Leveraging new software called inband migration tool (IMT), this feature supports allows for migrating a VNX VDM and file systems to a Dell EMC Unity NAS server. IMT was developed to handle the communication and data transfer between VNX and Dell EMC Unity. This enables ability for the migration to be transparent to host IO for NFS clients due to the stateless nature of the protocol. This feature allows for a full rollback to the VNX VDM without user data loss at any time before the migration is committed.

Some preparatory work is required on the VNX system, but creation, monitoring, and cutover are all managed on Dell EMC Unity. This feature is fully supported through Unisphere, as well as UEMCLI and REST API. Due to the Unified nature of Dell EMC Unity, only one remote system connection is required to leverage both the Native File Import and Native Block Import features.

For more information on Dell EMC Unity Native File Import, reference *the Dell EMC Unity: Migration Technologies* white paper and the *VNX Series Data Import to a Dell EMC Unity System User Guide* on [Dell EMC Support](#).

9 Conclusion

With the new file capabilities introduced in the Dell EMC Unity storage system, administrators gain access to a much more scalable, efficient, and high performing file system than previously available. These benefits are a result of the new 64-bit file system architecture, which also brings improvements in the areas of availability and recoverability. With a rich set of features, the Dell EMC Unity file system has both the power and flexibility to be leveraged for a wide array of traditional and transactional NAS use cases. Also important is the complete unification of file and block in a single platform. Dell EMC Unity bridges the gap between file and block with truly unified pools and features, allowing the core Dell EMC Unity feature set to equally benefit storage environments whether they are NAS, SAN, or a mixture of both.

A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

A.1 Related resources

Reference the following resources available on [Dell EMC Support](#):

- Dell EMC Unity: Best Practices Guide
- Dell EMC Unity: Cloud Tiering Appliance (CTA)
- Dell EMC Unity: Compression
- Dell EMC Unity: Compression for File
- Dell EMC Unity: Data at Rest Encryption
- Dell EMC Unity: Data Integrity
- Dell EMC Unity: Data Reduction
- Dell EMC Unity: DR Access and Testing
- Dell EMC Unity: Dynamic Pools
- Dell EMC Unity: FAST Technology Overview
- Dell EMC Unity: File-Level Retention (FLR)
- Dell EMC Unity: High Availability
- Dell EMC Unity: Introduction to the Platform
- Dell EMC Unity XT: Introduction to the Platform
- Dell EMC Unity: MetroSync
- Dell EMC Unity: MetroSync and Home Directories
- Dell EMC Unity: MetroSync and VMware vSphere NFS Datastores
- Dell EMC Unity: Migration Technologies
- Dell EMC Unity: OpenStack Best Practices for Ocata Release
- Dell EMC Unity: Performance Metrics
- Dell EMC Unity: Replication Technologies
- Dell EMC Unity: Snapshots and Thin Clones
- Dell EMC Unity: Operating Environment (OE) Overview
- Dell EMC Unity: Unisphere Overview
- Dell EMC Unity: Virtualization Integration
- Dell EMC UnityVSA
- Dell EMC Unity Cloud Edition with VMware Cloud on AWS
- Dell EMC Unity Data Reduction Analysis
- Dell EMC Unity: Migrating to Dell EMC Unity with SAN Copy
- Dell EMC Unity Storage with Microsoft Hyper-V
- Dell EMC Unity Storage with Microsoft SQL Server
- Dell EMC Unity Storage with Microsoft Exchange Server
- Dell EMC Unity Storage with VMware vSphere
- Dell EMC Unity Storage with Oracle Databases
- Dell EMC Unity 350F Storage with VMware Horizon View VDI
- Dell EMC Unity: 3,000 VMware Horizon Linked Clone VDI Users
- Dell EMC Storage with VMware Cloud Foundation