

Dell EMC Unity™ Family

Version 5.x

Configuring NFS File Sharing

H16959

REV 04

Copyright © 2018-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published June 2019

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Figures		7
Tables		9
Part 1	Basic functionality	11
Chapter 1	Overview	13
	Unity NFS support.....	14
	Unisphere storage provisioning.....	14
	Thin provisioning best practices.....	15
	Planning considerations.....	16
	Related features and functionality information.....	16
Chapter 2	Configuring NAS servers	19
	About secure NFS.....	20
	Create a NAS server for UNIX-only file sharing (NFS).....	20
	Configure NAS server sharing protocols and FTP/SFTP settings.....	22
	Configure a NAS server Unix Directory Service.....	23
	Configure Kerberos with a custom realm.....	26
	Change NAS server properties.....	28
	Change NAS server Unix credential settings.....	29
	View the active LDAPS CA certificate for a NAS server.....	30
	Upload an LDAPS CA certificate for a NAS server.....	30
	NDMP settings.....	30
Chapter 3	Configuring file systems	31
	Create a file system.....	32
	Change file system properties.....	32
	About Events Publishing.....	33
	Create Events Publishing notifications.....	34
	Change Events Publishing settings.....	35
	About automatic file system shrink and extend.....	35
	About manual file storage resource shrink and extend.....	36
	Manually shrink or extend the size of a file storage resource.....	36
Chapter 4	Configuring file system shares	39
	Share local paths and export paths.....	40
	Create an NFS share.....	40
	Change NFS share properties.....	41
Chapter 5	Performance metrics for NFS	43
	View historical performance metrics	44
	View real-time performance metrics.....	44
	File System Client Bandwidth.....	44
	File System Client Response Time.....	45

	File System Client I/O Size.....	45
	File System Client IOPS.....	45
	System - Client File System Bandwidth.....	45
	System - Client File System Response Time.....	46
	System - Client File System I/O Size.....	46
	System - Client File System IOPS.....	46
	System - NFS Bandwidth.....	46
	System - NFS I/O Size.....	47
	System - NFS IOPS.....	47
	System - NFS Response Time.....	47
	File System Bandwidth.....	47
	File System I/O Size.....	48
	File System IOPS.....	48
	System - File System Bandwidth.....	48
	System - File System I/O Size.....	48
	System - File System IOPS.....	49
	Tenant Bandwidth.....	49
Part 2	Advanced functionality	51
Chapter 6	Managing quotas	53
	About file system quotas.....	54
	Recommended approach for configuring quotas.....	55
	Quota policies.....	55
	Enable or disable the enforcement of user quotas on a quota tree.....	56
	Enable or disable the enforcement of user quotas on a file system.....	56
	Create a user quota on a file system.....	57
	Create a quota tree on a file system.....	57
	Create a user quota on a quota tree.....	58
	View file system storage space usage by user.....	58
	View quota tree storage space usage.....	58
	Change quota properties for a file system.....	58
	Change properties for a quota tree.....	59
	Change the quota policy for a file system.....	60
Chapter 7	Configure IP routes	61
	About NAS server routing.....	62
	NAS server interfaces.....	64
	Preferred interfaces for NAS servers.....	64
	IP Packet reflect functionality for NAS server interfaces.....	65
	Manage NAS server network interfaces and default routes.....	65
	Manage NAS server routes for responding to client requests.....	66
	Manage NAS server routes for external service requests.....	66
	Enable or disable IP packet reflect for a NAS server.....	67
	Verify NAS server routes.....	67
Chapter 8	Configuring IP multi-tenancy	69
	About IP multi-tenancy.....	70
	Configuring IP multi-tenancy.....	70
	Add a tenant.....	71
	Change tenant properties.....	71
	Configure file replication for a tenant	72

Chapter 9	Troubleshooting an NFS configuration	73
	Service commands for troubleshooting NFS issues in Unity.....	74

CONTENTS

FIGURES

1	Difference between thick and thin provisioning.....	15
---	---	----

FIGURES

TABLES

1	Differences between thick and thin provisioning.....	14
2	LDAP authentication.....	25
3	NAS server Unix credential settings.....	29
4	Event descriptions.....	34
5	Unity components for tenant T1.....	71
6	Unity components for tenant T2.....	71

TABLES

PART 1

Basic functionality

[Chapter 1, "Overview"](#)

[Chapter 2, "Configuring NAS servers"](#)

[Chapter 3, "Configuring file systems"](#)

[Chapter 4, "Configuring file system shares"](#)

[Chapter 5, "Performance metrics for NFS"](#)

CHAPTER 1

Overview

- [Unity NFS support](#)..... 14
- [Unisphere storage provisioning](#)..... 14
- [Thin provisioning best practices](#)..... 15
- [Planning considerations](#)..... 16
- [Related features and functionality information](#)..... 16

Unity NFS support

All Unity releases support NFSv3 and NFSv4. Unity also support secure NFS with Kerberos, for strong authentication. While Unity supports the majority of the NFSv4 and v4.1 functionality described in the relevant RFCs, directory delegation and pNFS are not supported.

NFS support is enabled on a NAS server during or after creation, allowing you to create NFS-enabled file systems on that NAS server.

Unisphere storage provisioning

Storage provisioning is the process of allocating available drive capacity to meet the capacity, performance, and availability requirements of hosts and applications. When you provision storage with Unisphere, you create storage resources to which hosts and applications can connect in order to access storage.

When you provision a storage resource in Unisphere, the system uses thin provisioning by default. This type of provisioning can improve storage efficiency while reducing the time and effort required for monitoring and rebalancing existing pool resources. Organizations can purchase less storage capacity up front, and increase available drive capacity (by adding drives) on an on-demand basis, and according to actual storage usage, instead of basing drive requirements in the requests or predictions of connected hosts. Thin provisioning allows multiple storage resources to subscribe to common storage capacity within a pool, while the system allocates only a portion of the physical capacity requested by each storage resource. The remaining storage is available for other storage resources to use.

Note

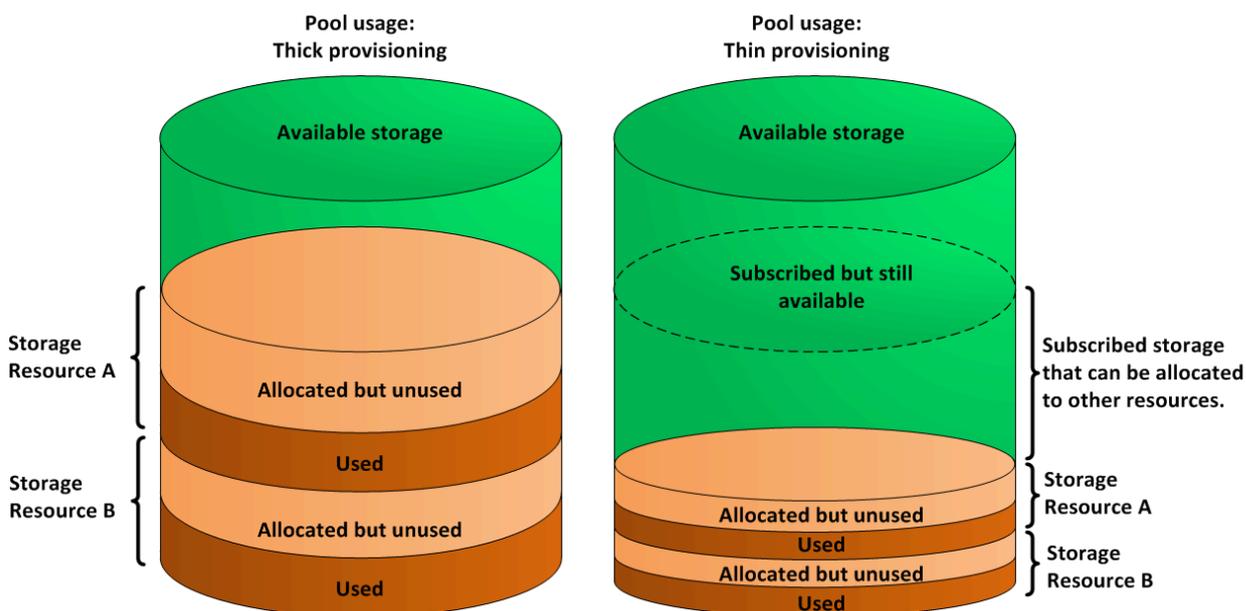
All storage resources require some amount of metadata from the pools where the storage resources were provisioned.

Thick and thin provisioning

The following table describes the differences between thick and thin provisioning:

Table 1 Differences between thick and thin provisioning

Provisioning type	Description
Thick provisioning	The amount of storage requested for a storage resource is exclusively allocated for it. This storage is reserved, and the unused portion cannot be used or distributed among other storage resources associated with the same pool.
Thin provisioning	The amount of storage requested for a storage resource is not immediately allocated for it. Instead, the system allocates an initial quantity of storage to the storage resource. When the amount of storage consumed within the storage resource approaches the limit of the current allocation, the system allocates additional storage to the storage resource from the pool. Thin provisioning is required for data reduction.

Figure 1 Difference between thick and thin provisioning**Creating a thin storage resource**

When you create a thin storage resource, you specify a target size for the resource. The size represents the maximum capacity to which the storage resource can grow without being increased by an administrator. The system reserves only a portion of the requested size, called the initial allocation. The requested size of the storage resource represents a subscribed quantity. Additional storage is allocated on-demand.

When a host or application uses approximately 75% of its initial allocation, an additional incremental quantity of storage is automatically allocated to the storage resource. The incremental allocation process continues until the quantity of storage allocated for the storage resource reaches the limit determined by its target size.

Note

A storage resource may appear full when data copied or written to the storage resource is greater than the space available at that time. When this occurs, the system begins to automatically extend the storage space and accommodate the write operation. As long as there is enough extension space available, this operation will complete successfully.

Pool subscription levels

Because storage resources can subscribe to more storage than is actually available to them, pools can be over-provisioned to support more storage capacity than they actually possess. The system automatically generates notification messages when total pool usage reaches 85% of the pool's physical capacity. (You can customize this threshold.)

Thin provisioning best practices

The following general rules can help determine the best environments in which to use thin provisioning:

- Thin provisioning provides the benefit of space efficiency. It is recommended that you choose thin provisioning for a storage resource (selected by default), unless absolute and predictable performance is a higher requirement than space efficiency. In some workload environments, performance can actually improve with thin provisioning.
- Environments that can benefit from thin provisioning include:
 - Document repositories with rapidly rising capacity requirements. These repositories can benefit greatly from the improved capacity utilization offered by thin provisioning, provided their environments meet the previously outlined criteria.
 - Software development and source code repositories. These repositories are well-suited to thin provisioning, because their environments can usually tolerate some level of performance variability.
- Thin provisioning works best in file system environments where files are not frequently deleted. Many file systems do not efficiently reuse the space associated with deleted files, which can result in an allocated but unused space in the thin-provisioned file system.
- Consider the space consumption characteristics of databases before using thin provisioning. Some databases pre-allocate the storage space for data before writing to it. This space is allocated within a thin-provisioned storage resource, and this can reduce the capacity utilization within the pool. For more information, consult your database vendor documentation.

Advantages of thin and standard provisioning

Thin provisioning provides the following advantages:

- Provides the most efficient allocation of storage capacity based on usage.
- Promotes ease of use in setting up and managing pool capacity.
- Minimizes the host impact of adding pool resources based on host storage usage.
- Optimizes storage usage in situations where space consumption is difficult to forecast.

Planning considerations

The following table summarizes the tasks to perform before you start configuring NFS on your Unity system. For more information on performing these tasks, see the Unity online help.

1. Optionally configure at least one NTP server on the storage system to synchronize the date and time. It is recommended that you set up a minimum of two NTP servers per domain to avoid a single point of failure. This step is mandatory if you are using secure NFS.
2. Optionally configure VLANs and tenants if you plan to implement multi-tenancy.
3. Optionally configure a Unix Directory Service. This step is mandatory if you are using secure NFS, unless you use local files.
4. Optionally configure one or more DNS servers. This step is mandatory if you are using secure NFS.

Related features and functionality information

Specific information related to the features and functionality described in this document is included in the following for Unity:

- Unisphere Online Help
- *Configuring Hosts to Access NFS File Systems*
- *Configuring Replication*
- *Unisphere Command Line Interface User Guide*
- *Service Commands Technical Notes*

The complete set of customer publications is available on the Online Support website at <http://Support.EMC.com>. After logging in to the website, click the **Support by Product** page and specify **Dell EMC Unity Family** to locate information for the specific feature required.

CHAPTER 2

Configuring NAS servers

- [About secure NFS](#).....20
- [Create a NAS server for UNIX-only file sharing \(NFS\)](#).....20
- [Configure NAS server sharing protocols and FTP/SFTP settings](#)..... 22
- [Configure a NAS server Unix Directory Service](#)..... 23
- [Configure Kerberos with a custom realm](#)..... 26
- [Change NAS server properties](#).....28
- [Change NAS server Unix credential settings](#).....29
- [View the active LDAPS CA certificate for a NAS server](#)30
- [Upload an LDAPS CA certificate for a NAS server](#) 30
- [NDMP settings](#)..... 30

About secure NFS

You can configure secure NFS when you create or modify a NAS server that supports Unix shares. Secure NFS provides Kerberos-based user authentication, which can provide network data integrity and network data privacy.

Kerberos is a distributed authentication service designed to provide strong authentication with secret-key cryptography. It works on the basis of "tickets" that allow nodes communicating over a non-secure network to prove their identity in a secure manner. When configured to act as a secure NFS server, the NAS server uses the RPCSEC_GSS security framework and Kerberos authentication protocol to verify users and services.

Security options

Secure NFS supports the following security options:

- krb5: Kerberos authentication.
- krb5i: Kerberos authentication and data integrity by adding a signature to each NFS packet transmitted over the network.
- krb5p: Kerberos authentication, data integrity, and data privacy by encrypting the data before sending it over the network. Data encryption requires additional resources for system processing and can lead to slower performance.

In a secure NFS environment, user access to NFS file systems is granted based on Kerberos principal names. However, access control to shares within a file system is based on the Unix UID and GID, or on ACLs.

Note

Secure NFS supports NFS credentials with more than 16 groups. This is equivalent to the extended Unix credentials option.

Configuring secure NFS

To configure secure-NFS for a NAS server that supports NFS only, configure a custom realm to point to any type of Kerberos realm (AD, MIT, Heimdal). You must upload the keytab file to the NAS server being defined.

Create a NAS server for UNIX-only file sharing (NFS)

Before you begin

Obtain the following information:

- (Optional) Name of the tenant to associate with the NAS server.
- Name of the pool to store the NAS server's metadata.
- Storage Processor (SP) on which the NAS server will run.
- IP address information for the NAS server.
- VLAN ID, if the switch port supports VLAN tagging. If you associate a tenant with the NAS server, you must choose a VLAN ID.
- (Optional) UNIX Directory Service (UDS) information for NIS or LDAP, or local files. This can be used to resolve hosts defined on NFS share access lists.
- (Optional) DNS server information. This can also be used to resolve hosts defined on NFS share access lists.

- (Optional) Replication information.

It is recommended that you balance the number of NAS servers on both SPs.

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the **Add** icon.
3. On the **General** and **Interface** pages, specify the relevant settings. Note the following:
 - On the **General** page, the **Server name** identifies the NAS server. It is not a network name.
 - Optionally select a tenant to associate with the NAS server.

Note

Once you create a NAS server that has an associated tenant, you cannot change this association.

-
- On the **Interface** page, optionally select a VLAN. If you selected a tenant on the **General** page, you must select a VLAN. The list of VLANs represent the VLANs associated with the selected tenant.
 4. On the **Sharing Protocols** page:
 - Select **Linux/Unix shares (NFS)**.
 - Select whether to enable NFSv3, NFSv4, or both.
 - Optionally enable support for Virtual Volumes (VVols).
 - Optionally click **Configure secure NFS** to enable secure NFS with Kerberos. When you enable secure NFS for a NAS server that supports Unix-only file sharing, you must configure a custom Kerberos realm.
 5. On the **Unix Directory Service** page, configure one of the following directory services (optional unless you are configuring secure NFS):
 - Local files
 - NIS
 - LDAP
 - Local files and NIS
 - Local files and LDAP

If you configure local files with NIS or LDAP, the system queries the local files first. You can configure LDAP to use anonymous, simple, and Kerberos authentication. You can also configure LDAP with SSL (LDAP Secure) and can enforce the use of a Certificate Authority certificate for authentication.
 6. On the **DNS** page, optionally configure DNS for the NAS server.
 7. On the **Replication** page, optionally select a replication mode and Recovery Point Objective (RPO) for the NAS server.

Configure NAS server sharing protocols and FTP/SFTP settings

You can configure NFS support when you create a NAS server or change its properties. You can configure FTP/SFTP support for an existing NAS server only.

If you are creating a NAS server, access the NAS server sharing protocol options from the **Sharing Protocols** window in the **Create a NAS server** wizard.

If you are changing NAS server properties, follow these steps to access the NAS server sharing protocol and FTP options:

1. Under **Storage**, select **File > NAS Servers**.
2. Select the relevant NAS server, and then select the **Edit** icon.
3. Select the **Sharing Protocols** tab.

NFS support

If you are changing NAS server properties, select the **NFS** sub-tab on the **Sharing Protocols** tab.

Task	Description
Enable or disable the NAS server's ability to serve files using the NFS protocol.	Select or clear the Enable Linux/Unix shares (NFS Server) option. <ul style="list-style-type: none"> • To enable NFSv3, select Enable NFSv3 (when creating a NAS server) or NFSv3 enabled (when editing NAS server properties). • To enable NFSv4, select Enable NFSv4 (when creating a NAS server) or NFSv4 enabled (when editing NAS server properties).
Enable or disable the NAS server's ability to serve VVols.	Select or clear Enable VVols . If you enable VVols, you must select the IP address for the VMware protocol endpoint.
Enable or disable support for secure NFS.	Select Show advanced , and then select or clear Enable Secure NFS (with Kerberos) .
Configure secure NFS using a custom realm	See Configure Kerberos with a custom realm on page 26.

FTP/SFTP support

You can configure FTP or FTP over SSH (SFTP) settings for an existing NAS server only. Select the **FTP** sub-tab on the **Sharing Protocols** tab.

Task	Description
Enable or disable the NAS server's ability to share files using the FTP protocol.	Select or clear Enable FTP . If this option is selected, optionally click the other options to customize user authentication, user home directory, and message settings.

Task	Description
Enable or disable the NAS server's ability to share files using the SFTP protocol.	Select or clear Enable SFTP . If this option is selected, optionally click the other options to customize user authentication, user home directory, and message settings.

FTP access can be authenticated using the same methods as NFS. Once authentication is complete, access is the same as NFS for security and permission purposes. If the format is anything other than `domain@user` or `domain\user`, NFS authentication is used. NFS authentication uses local files, LDAP, NIS, or local files with LDAP or NIS.

To use local files for FTP access, the `passwd` file must include an encrypted password for the users. This password is used for FTP access only. The `passwd` file uses the same format and syntax as a standard Unix system, so you can leverage this to generate the local `passwd` file. On a Unix system, use `useradd` to add a new user and `passwd` 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 the `/etc/passwd` file to the NAS server.

Configure a NAS server Unix Directory Service

There are three ways to configure identity lookups:

- [Use local files](#), alone or with a UDS.
- [Configure a Unix Directory Service \(UDS\) using NIS](#).
- [Configure a UDS using LDAP](#).

Note

If you configure local files with a UDS, the storage system queries the local files first.

If you are creating a new NAS server, use the **Unix Directory Service** window in the **Create a NAS server** wizard to configure identity lookups.

If you are configuring a UDS for an existing NAS server, access the **Naming Services** tab to access the identity lookup options:

1. Under **Storage**, select **File > NAS Servers**.
2. Select a NAS server, and then select the **Edit** icon.
3. Select the **Naming Services** tab.

Using local files

To enable the use of local files for directory services when you are creating a NAS server:

1. From the **Unix Directory Service** window in the **Create a NAS server** wizard, select **Enable a Unix Directory service using Local Files**.
2. Create the password file for the UDS. To view the template for this file, select **Open a Passwd File Template**.
3. Select **Upload Passwd File** to upload the password file to the NAS server.

After you create the NAS server, you can upload additional local files as specified below.

To enable the use of local files for directory services for an existing NAS server:

1. From the **Naming Services** tab, select the **Local Files** sub-tab.
2. Select **Enable a Unix Directory service using Local Files**.
3. For each type of local file, select **Retrieve current <file-type> file** to download the current file. If there is no file on the storage system, the system downloads a file template.
4. Make the necessary changes to the file.
5. Select **Upload New <file-type> File** to upload the file.

To troubleshoot issues with configuring local files, ensure that:

- The file is created with the proper syntax. (Six colons are required for each line). Reference the template for more details about the syntax and examples.
- Each user has a unique name and UID.

Configuring a Unix Directory Service using NIS

To configure a UDS using NIS when you are creating a NAS server:

1. From the **Naming Services** tab, select the **LDAP/NIS** sub-tab.
2. In the **Enable Unix Directory service** field, select **NIS**.
3. Enter an NIS domain and add up to three IP addresses for the NIS servers.

To configure a UDS using NIS for an existing NAS server:

1. From the **Naming Services** tab, select the **LDAP/NIS** sub-tab.
2. In the **Enable Unix Directory service** field, select **NIS**.
3. Enter an NIS domain and add up to three IP addresses for the NIS servers.

To troubleshoot issues with configuring a UDS using NIS, ensure that the NIS server domain and server IP addresses you enter are correct.

Configure a UDS using LDAP

LDAP must adhere to the IDMU, RFC2307, or RFC2307bis schemas. Some examples include AD LDAP with IDMU, iPlanet, and OpenLDAP. The LDAP server must be configured properly to provide UIDs for each user. For example, on IDMU, the administrator must go in to the properties of each user and add a UID to the UNIX Attributes tab.

To configure a UDS using LDAP when you are creating a NAS server:

1. From the **Naming Services** tab, select the **LDAP/NIS** sub-tab.
2. In the **Enable Unix Directory service** field, select **LDAP**.
3. Select how the NAS server will obtain LDAP server IPs:
 - If you leave the default option, the NAS server will use DNS service discovery to obtain LDAP server IP addresses automatically. For this discovery process to work, the DNS server must contain pointers to the LDAP servers, and the LDAP servers must share the same authentication settings.
 - To manually enter the IP addresses of LDAP servers, select **Configure LDAP server IPs manually**, enter each IP address, and click **Add**.
4. Configure the LDAP authentication as described in [Table 2](#) on page 25.

To configure a UDS using LDAP for an existing NAS server:

1. From the **Naming Services** tab, select the **LDAP/NIS** sub-tab.
2. In the **Enable Unix Directory service** field, select **LDAP**
3. Configure the LDAP authentication as described in [Table 2](#) on page 25.

Note

By default, LDAP uses port 389, and LDAPS (LDAP over SSL) uses port 636.

Table 2 LDAP authentication

Option	Considerations
LDAP with Anonymous or Simple authentication	<p>For Anonymous Authentication, add the LDAP servers and specify the port number used by the LDAP servers, the Base DN, and the Profile DN for the iPlanet/OpenLDAP server.</p> <p>For Simple Authentication, add the LDAP servers and specify the following:</p> <ul style="list-style-type: none"> • If using AD, LDAP/IDMU: <ul style="list-style-type: none"> ▪ Port number used by the LDAP servers. ▪ User account in LDAP notation format; for example, cn=administrator,cn=users,dc=svt,dc=lab,dc=com. ▪ User account password. ▪ Base DN, which is the same as the Fully Qualified Domain Name (for example, svt.lab.com). • If using the iPlanet/OpenLDAP server: <ul style="list-style-type: none"> ▪ User account in LDAP notation format; for example, cn=administrator,cn=users,dc=svt,dc=lab,dc=com. ▪ Password. ▪ Base DN. For example, if using svt.lab.com, the Base DN would be DC=svt,DC=lab,DC=com. ▪ Profile DN for the iPlanet/OpenLDAP server.
LDAP with Kerberos authentication	<p>To configure Kerberos, configure a custom realm to point to any type of Kerberos realm (Windows, MIT, 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 of the SMB server is not used when you choose this option.</p> <hr/> <p>Note</p> <p>If you use NFS secure with a custom realm, you have to upload a keytab file.</p>

To troubleshoot issues with configuring a UDS using LDAP, ensure that:

- The LDAP configuration adheres to one of the supported schemas, as described earlier in this topic.
- All of the containers specified in the `ldap.conf` file point to containers that are valid and exist.
- Each LDAP user is configured with a unique UID.

You can also use the `-ldap` option of the `svc_nas service` command to troubleshoot LDAP issues. This command can display advanced diagnostics for the connection to the LDAP server and can run a user name resolution to ensure that the LDAP settings are correct. For more information, see the *Service Commands Technical Notes*, which is available from the [UnityOE Features Info Hub](#).

Configure Kerberos with a custom realm

This method of configuring Kerberos lets you configure any kind of KDC (MIT/Heidmal or AD). Use this method when you do not have an SMB server domain configured on the NAS server or if you want to use a different Kerberos realm than the one configured for the SMB server.

If you are configuring Kerberos for secure NFS, be aware of the following:

- Using LDAPS or LDAP with Kerberos is recommended for increased security.
- A DNS server must be configured at the NAS-server level. All members of the Kerberos realm, including the KDC, NFS server, and NFS clients, must be registered in the DNS server. Some applications, such as VMware, might also require reverse DNS lookup.
- The NFS client's hostname FQDN and NAS server FQDN must be registered in the DNS server. Clients and servers must be able to resolve any member of the Kerberos realm's FQDNs to an IP address.
- The FQDN part of the NFS client's SPN must be registered in the DNS server.

Note

To configure Kerberos, the storage system must have a configured NTP server. Kerberos relies on the correct time synchronization between the KDC, servers, and client on the network.

Before using Unisphere

To use a Windows-based KDC without using the SMB server account on the NAS server, follow these steps before configuring Kerberos in Unisphere. The steps assume you want to use `myrealm.windows.dellemc.com` as the FQDN for the NFS server.

1. Create account `myrealm` for the NAS server in the Active Directory (AD) of the windows domain `windows.dellemc.com`.
2. Register the service SPN on the computer account you created:

```
C:\setspn -S nfs/myrealm.windows.dellemc.com myrealm
```

3. Verify that the SPN was created.

```
C:\setspn myrealm
```

4. Generate a keytab file for the SPN:

```
C:\ktpass -princ nfs/
myrealm.windows.dellemc.com@WINDOWS.DELLEM.COM -mapuser WINDOWS
\myrealm
-crypto ALL +rndpass -ptype KRB5_NT_PRINCIPAL -out
myrealm.windows.dellemc.com.keytab
```

To use a Unix-based KDC, follow these steps before configuring Kerberos in Unisphere. The steps assume you want to use `myrealm` in the Kerberos realm `linux.dellemc.com` as the hostname of the NFS server.

1. Run the `kadmin.local` tool.

2. Create the principals and their keys:

```
kadmin.local: addprinc -randkey nfs/myrealm.linux.dellemc.com
```

and/or

```
kadmin.local: addprinc -randkey nfs/myrealm
```

3. Put the key of the principal into the keytab file myrealm.linux.dellemc.fr:

```
kadmin.local: ktadd -k myrealm.linux.dellemc.com.keytab nfs/  
myrealm.linux.dellemc.fr
```

When creating a NAS server

To configure Kerberos with a custom realm when you create a NAS server, follow the steps in the Create a NAS Server wizard, while noting the following:

- If you are configuring Kerberos for secure NFS:
 1. On the **Sharing Protocols** window configure a NAS server that supports NFS or multiprotocol file sharing.
 2. Select **Configure secure NFS**.
 3. Select **Enable Secure NFS (with Kerberos) > Use custom realm**.
 4. Enter the name of the custom realm.
 5. Upload the keytab file to the NAS server's NFS server.
 6. On the **Unix Directory Service** window, add the LDAP servers, and specify the Kerberos principal, password, base DN, and optionally, profile DN.
 7. On the **DNS** window, configure DNS for the NAS server.
 8. Register all members of the Kerberos realm in the DNS server.
- If you are configuring Kerberos for LDAP or LDAP secure:
 1. On the **Sharing Protocols** window configure a NAS server that supports NFS or multiprotocol file sharing.
 2. On the **Unix Directory Service** window, add the LDAP servers and select **Kerberos** as the authentication method.
 3. Specify the principal, password for the principal, and base DN.
 4. On the **Kerberos** window, add the KDC servers, and optionally change the TCP port.
 5. On the **DNS** window, configure DNS for the NAS server.

When changing NAS server properties

- If you are configuring Kerberos for secure NFS:
 1. Make sure that DNS and a UDS are configured for the NAS server and that all members of the Kerberos realm are registered in the DNS server.
 2. On the **Security** tab, select the **Kerberos** sub-tab, and then select **Configure custom Kerberos** settings.
 3. Configure the custom Kerberos settings.
 4. Upload the keytab file to the NAS server's NFS server.
 5. On the **Sharing Protocols** tab, select the **NFS** sub-tab.

6. Select **Show advanced**, and specify the host name.
 7. Select **Enable Secure NFS (with Kerberos) > Use custom realm**.
- If you are configuring Kerberos for LDAP:
 1. Make sure that DNS and LDAP are configured for the NAS server and that all members of the Kerberos realm are registered in the DNS server.
 2. On the **Security** tab, select the **Kerberos** sub-tab, and then select **Configure custom Kerberos** settings.
 3. Configure the custom Kerberos settings.
 4. On the **Naming Services** tab, select the LDAP/NIS sub-tab, and select **Kerberos** as the LDAP authentication method.
 5. Select **Specify custom principal**.
 6. Specify the principal and password for the principal.

Troubleshooting Kerberos

You can use the `-kerberos` option of the `svc_nas` service command to troubleshoot Kerberos issues. For more information, see the *Service Commands Technical Notes*, which is available from the [UnityOE Features Info Hub](#).

Change NAS server properties

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the relevant NAS server, and then select the **Edit** icon.
3. On the **General** tab:
 - Change the NAS server name.
 - Select **SP Owner** to transition from one SP to the other SP for this NAS server. For example, you may want to do this if you have an overloaded SP, and want to reduce the load by moving the server to the other SP.
4. On the **Network** tab:
 - Select the **Interfaces & Routes** sub-tab to add, change, delete, or verify NAS server interfaces, enable or disable IP packet reflect for the NAS server, or change the NAS server's preferred interfaces. Select an interface, and then select **Show external routes for interfaces** to access the per-interface routing table, where you can add, change, or delete the selected interface's routes for responding to client requests.
 - Select the **Routes to External Services** sub-tab to add, change, or verify NAS server routes for external service requests, or to configure default gateways.
5. On the **Naming Services** tab, configure DNS and either configure the UNIX Directory Service (UDS) for the NAS server (LDAP or NIS) or use local files. Alternatively, you can use local files with a UDS. In this case, the system checks the local files first.
6. On the **Sharing Protocols** tab:
 - Select the **NFS** sub-tab to enable or disable support for NFS shares, VVols, NFSv3, NFSv4, and extended UNIX credentials. You can also configure secure NFS with Kerberos and change the credential cache retention period.

- Select the **FTP** sub-tab to enable or disable FTP or SFTP, or to change FTP or SFTP properties.
7. On the **Protection & Events** tab:
 - Select the **NDMP Backup** sub-tab to enable or disable NDMP, and to change the NDMP password.
 - Select the **DHSM** sub-tab to enable or disable Distributed Hierarchical Storage Management (DHSM) and to change the DHSM password.
 - Select the **Events Publishing** sub-tab to enable or disable Events Publishing, create or modify an event pool, and create or modify events policy settings.
 8. On the **Security** tab, select the **Kerberos** sub-tab to configure a custom Kerberos realm and to retrieve or upload the Kerberos keytab file.
 9. On the **Replication** tab, optionally select a replication mode and Recovery Point Objective (RPO) for the NAS server.

Change NAS server Unix credential settings

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the relevant NAS server from the list, and then select the **Edit** icon.
3. On the **Sharing Protocols** tab, select **Show advanced**.
4. Make the desired changes, as described in the following table.

Table 3 NAS server Unix credential settings

Task	Description
Extend the Unix credential to enable the storage system to obtain more than 16 group GIDs. Note With secure NFS, the Unix credential is always built by the NAS server, so this option does not apply.	Select or clear Enable extended Unix credentials . <ul style="list-style-type: none"> • If this field is selected, the NAS server uses the User ID (UID) to obtain the primary Group ID (GID) and all group GIDs to which it belongs. The NAS server obtains the GIDs from the local password file or UDS. • If this field is cleared, the Unix credential of the NFS request is directly extracted from the network information contained in the frame. This method has better performance, but it is limited to including up to only 16 group GIDs.
Specify a Unix credential cache retention period. This option can lead to better performance, because it reuses the Unix credential from the cache instead of building it for each request.	In the Credential cache retention field, enter a time period (in minutes) for which access credentials are retained in the cache. The default value is 15 minutes, minimum value is 1 minute, and maximum value is 1439 minutes.

View the active LDAPS CA certificate for a NAS server

This option is available for anonymous and simple LDAP authentication that uses SSL and enforces certification.

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the relevant NAS server from the list, and then select the **Edit** icon.
3. Select the **Naming Services** tab, and then select the **LDAP/NIS** sub-tab.
4. Click **Retrieve CA Certificate**.

Upload an LDAPS CA certificate for a NAS server

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the relevant NAS server, and click the **Edit** icon.
3. On the **Naming Services** tab, select the **LDAP/NIS** sub-tab.
4. Select **LDAP Secure (Use SSL)** and **Enforce Certification Authority (CA) Certificate**, if these options are not already selected. These options are available for Anonymous and Simple authentication.
5. Select **Upload CA Certificate**, locate the certificate to upload, locate the certificate, and click **Start Upload**.

NDMP settings

The Network Data Management Protocol (NDMP) provides a standard for backing up file servers on a network. NDMP allows centralized applications to back up file servers running on various platforms and platform versions. NDMP reduces network congestion by isolating control path traffic from data path traffic, which permits centrally managed and monitored local backup operations. Enabling NDMP for file system storage resources makes it possible to use third party NDMP products to back up and restore file system data.

You can enable NDMP by configuring NAS server settings.

CHAPTER 3

Configuring file systems

- [Create a file system](#)..... 32
- [Change file system properties](#)..... 32
- [About Events Publishing](#)..... 33
- [Create Events Publishing notifications](#)..... 34
- [Change Events Publishing settings](#)..... 35
- [About automatic file system shrink and extend](#)..... 35
- [About manual file storage resource shrink and extend](#)..... 36
- [Manually shrink or extend the size of a file storage resource](#)..... 36

Create a file system

Before you begin

Make sure there is a NAS server configured to support the NFS file system type, and that a pool exists with enough available storage space.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the **Add** icon.
3. On the **Protocol** window, select the **Linux/Unix Shares (NFS)**. Then select the associated NAS server.
4. Continue following the steps in the wizard while noting the following:
 - On the **Storage** page, the **Thin** checkbox is selected by default. If you do not want to create a thin file system, remove the checkmark from the **Thin** checkbox. Removing the checkmark also disables the **Data Reduction** option.
 - On the **Storage** page, select the **Data Reduction** checkbox to enable data reduction on the file system. Data reduction is applied on all new incoming writes. Data that already exists on the file system does not have data reduction applied. Data reduction can be enabled only on thin file systems that reside in All-Flash pools, and only for thin file systems created on Unity systems running OE version 4.2.x or later.
 - On the **Shares** page, optionally, configure the initial share for the file system.
 - You can configure host access and a snapshot schedule for the file system when you create the file system, or you can do this at a later time.

Change file system properties

If the associated NAS server is a replication destination, many configuration options cannot be changed.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the relevant file system, and then select the **Edit** icon.
3. On the **General** tab:
 - You can change the description of the file system and the file system size.
 - You can view the file system capacity, including used space and free space, on this tab.
 - If a thin file system was created on a Unity system running OE version 4.1, you can also change the minimum allocation size. You cannot reduce the storage size lower than the current allocation.
 - You can enable data reduction for thin file systems created on a Unity system running OE version 4.2.x or later. As data reduction is applied to all new incoming writes to a file system, data reduction statistics (such as data reduction ratio) display on the Properties page.

If data reduction is enabled and then subsequently disabled, existing data in the file system will remain as is, but newly-written data will not have data reduction applied.

- If data reduction is enabled, you can also enable Advanced Deduplication, which provides the ability to reduce the amount of data storage needed by eliminating redundant data from the system. Once enabled, all incoming writes to the system will have advanced deduplication applied. Advanced deduplication is available only on:
 - Dynamic or Traditional pools in Unity 380F, 480F, 680F, and 880F systems
 - Dynamic pools in Unity All-Flash 450F, 550F, and 650F systems
 - All-Flash pools in Unity Hybrid 380, 480, 680, and 880 systems
 - You can change capacity alarm threshold settings for when Info, Warning, and Error alert messages are generated.
4. On the **Snapshots** tab, manage the file system's snapshots or configure a snapshot schedule for the file system.
 5. On the **FAST VP** tab, change the file system tiering policy and view the data distribution per tier.
 6. On the **Advanced** tab, optionally enable Events Publishing for a file system.
 7. On the **Quota** tab, configure or change settings for file system quotas and quota trees.
 8. On the **Replication** tab, configure or change the file system replication settings.

Note

Replication can be set on the file-system level only if the replication session already exists for the NAS server where the file system resides.

9. On the **FLR** tab (FLR-enabled file systems only), optionally modify the retention periods, enable auto-lock of new files, set an auto-lock policy interval, or enable automatic deletion of files once the retention period expires.

Note

If the file system is a replication destination, FLR properties cannot be modified.

About Events Publishing

Events Publishing allows third-party applications to register to receive event notification and context from the storage system when accessing file systems by using the NFS protocols. The Events Publishing agent delivers to the application both event notification and associated context in one message. Context may consist of file metadata or directory metadata that is needed to decide business policy.

You must define at least one event option (pre-, post-, or post-error event) when Events Publishing is enabled.

- Pre-event notifications are sent before processing an NFS client request.
- Post-event notifications are sent after a successful NFS client request.

- Post-error event notifications are sent after a failed NFS client request.

Table 4 Event descriptions

Value	Definition	Protocol
OpenFileNoAccess	Sends a notification when a file is opened for a change other than read or write access (for example, read or write attributes on the file).	NFSv4
OpenFileRead	Sends a notification when a file is opened for read access.	NFSv4
OpenFileReadOffline	Sends a notification when an offline file is opened for read access.	NFSv4
OpenFileWrite	Sends a notification when a file is opened for write access.	NFSv4
OpenFileWriteOffline	Sends a notification when an offline file is opened for write access.	NFSv4
FileRead	Sends a notification when a file read is received over NFS.	NFSv3 or NFSv4
FileWrite	Sends a notification when a file write is received over NFS.	NFSv3 or NFSv4
CreateFile	Sends a notification when a file is created.	NFSv3 or NFSv4
CreateDir	Sends a notification when a directory is created.	NFSv3 or NFSv4
DeleteFile	Sends a notification when a file is deleted.	NFSv3 or NFSv4
DeleteDir	Sends a notification when a directory is deleted.	NFSv3 or NFSv4
CloseModified	Sends a notification when a file is changed before closing.	NFSv4
CloseUnmodified	Sends a notification when a file is not changed before closing.	NFSv4
RenameFile	Sends a notification when a file is renamed.	NFSv3 or NFSv4
RenameDir	Sends a notification when a directory is renamed.	NFSv3 or NFSv4
SetSecFile	Sends a notification when a file security change is received over NFS.	NFSv3 or NFSv4
SetSecDir	Sends a notification when a directory security change is received over NFS.	NFSv3 or NFSv4

Create Events Publishing notifications

Before you begin

Before you can set up Events Publishing for a NAS server:

- You cannot enable Events Publishing for a NAS server that is acting as a replication destination.
- At least one file system must exist for the NAS server.
- You must obtain the IP addresses of the CEPA servers.
- Ensure that NFS protocol events notifications have been enabled on the **File Systems Properties Advanced** window.

Procedure

1. Under **Storage**, select **File > NAS Servers**.

2. Select the relevant NAS Server, and then select the **Edit** icon.
3. On the **Protection & Events** tab, select the **Events Publishing** sub-tab.
4. Select the **Enable Common Event Publishing** checkbox.
5. On the **New Event Pool** window, specify the required items. You must configure at least one event from one of the available categories (pre-event, post-event, or post-error event).
6. Click **Configure**.
7. Optionally, select **Show policy settings** to configure pre-events and post-events failure policies.
8. Optionally, select **Show advanced settings** to configure CEPA server options.
9. Click **Apply** after you finish configuring events.

Change Events Publishing settings

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the relevant NAS Server, and then select the **Edit** icon.
3. On the **Protection & Events** tab, select the **Events Publishing** sub-tab.
4. To change event pool settings, select the pool and then the **Edit** icon. Update any necessary information.
5. To change event policy or connection settings, select the pool and then the relevant **Show settings** link. Update any necessary information.

About automatic file system shrink and extend

The system automatically shrinks and extends a thin file system based on capacity needs.

Thin-provisioned file systems are automatically extended or shrunk by the system when certain conditions are met. Automatic extend prevents the file system from running out of space. Automatic shrink improves space allocation by releasing any unused space back to the storage pool. The automatic shrink and automatic extend operations are based on a high water mark (HWM) for auto-extend and a low water mark (LWM) for auto-shrink.

For file systems smaller than 2.5 terabytes (TiB) in size:

- The file system is automatically extended when the used space exceeds and sustains over 75% of the allocated space. This is the fixed high water mark (HWM) for file systems less than 2.5 TiBs.
- The file system automatically shrinks and returns space to the pool when the used space is 70% less than the allocated space. This is the fixed low water mark (LWM) for file systems less than 2.5 TiBs.

For file systems larger than 2.5 TiBs in size, the high and low watermarks will be dynamic and operate based on the following:

- Auto-extend will wait until nearly all of the allocated space capacity is used before extending file systems larger than 2.5 TiBs.
- Auto-shrink will not require a large amount of capacity to be freed back to the pool as part of the shrink operation.

For larger file systems greater than 2.5 TiBs, the maximum extend size is 1 TiB. This helps avoid over-allocation of space from the pool to that file system that may not be immediately used.

You can set a minimum allocated size for a thin file system; automatic and manual shrink operations will not be able to reduce the size of the file system below this minimum. The default minimum allocated size for a thin file system is 3G.

About manual file storage resource shrink and extend

You can manually extend or shrink file systems.

File resource shrink

The shrink operation reduces the space the file resource uses from the pool, allowing the pool to reclaim the free, unused space from the target file resource.

For thick-provisioned file resources, you can shrink the size of the resource and return unused space to the pool. For example, if a thick file system is shrunk from a size of 1 TB down to 500 GB:

- The amount of used space for the resource remains the same.
- The free space for the resource is reduced by 500 GB.
- The total pool free space is increased by slightly less than 500 GB.
- The pool size used for the resource is reduced to approximately 500 GB.

The system displays a message indicating exactly how much space will be reclaimed by the pool as a result of the shrink operation.

For thin-provisioned file resources, you can manually shrink the size of a file resource to a size that is equal to or less than the allocated size.

Note

For Unity systems running OE version 4.1.x, the minimum size of a thin storage resource is 3 GB. You cannot shrink a thin file resource below the size used. For Unity systems running OE version 4.2 or later, the thin file storage resource minimum allocated size option is not supported.

File resource extend

The manual extend operation does the following for thin- and thick-provisioned file resources:

- For thin-provisioned file resources, increases the visible (virtual) size of the resource without increasing the actual size allocated to the resource from the pool.
- For thick-provisioned file resources, increases the actual space allocated to the resource from the pool.

Note

You cannot extend a thick file resource beyond the total pool free size.

Manually shrink or extend the size of a file storage resource

The ability to manually shrink or extend the size of a storage resource applies to file systems. A manual shrink allows the pool to reclaim space, while a manual extend allocates more space to the storage resource.

Note

You can cancel a manual shrink operation, but the progress made prior to cancellation will not be reverted.

Procedure

1. Select a storage resource, and then click the **Edit** icon.
 2. In the **Size** field, enter the new reduced (shrink) or increased (extend) size of the storage resource.
-

Note

For Unity systems running OE version 4.1.x, the minimum size of a storage resource is 3 GB. You cannot shrink below the size used or extend beyond the total pool free size.

CHAPTER 4

Configuring file system shares

- [Share local paths and export paths](#)..... 40
- [Create an NFS share](#)..... 40
- [Change NFS share properties](#)..... 41

Share local paths and export paths

The following table describes the path settings for shares:

Setting	Description
Local path	<p>The path to the file system storage resource on the storage system. This path specifies the unique location of the share on the storage system.</p> <p>NFS shares</p> <ul style="list-style-type: none"> • Each NFS share must have a unique local path. Unisphere automatically assigns this path to the initial share created within a new file system. The local path name is based on the file system name. • Before you can create additional shares within an NFS file system, you must create a directory to share from a Linux/UNIX host that is connected to the file system. Then, you can create a share from Unisphere and set access permissions accordingly.
Export path	<p>The path used by the host to connect to the share. Unisphere creates the share export path based on the name of the share and the name of the file system where it resides. Hosts use either the file name or the export path to mount or map to the share from a network host.</p> <p>This behavior is enabled by using NFS aliases for shares.</p>

Create an NFS share

Before you begin

The file system or snapshot you choose as the share's source must be associated with a NAS server that supports the NFS protocol.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system for which you want to add a share, and then select **More Actions > Create an NFS share (NFS export)**.
3. On the **File System** page, specify whether the share is for the selected file system or for a snapshot of the selected file system.
4. On the **Name & Path** page, enter the relevant information, noting the following:
 - The value specified in the **Share Name** field, along with the NAS server name, constitutes the alias by which hosts can access the share.
 - Share names must be unique at the NAS server level per protocol. However, you can specify the same name for an SMB and NFS share.
 - **Local Path** must correspond to an existing folder name within the file system that was created from the host-side.

Note

A given file system path can only be shared once using the NFS protocol.

- By default, users can set bit `s` in the execute portion of the owner or group permissions of a file. Users can then set the `setuid` and `setgid` Unix permission bits. This allows users to run the executable with the privileges of the file's owner (such as root). De-select **Allow SUID** if you do not want users to have this ability.
 - Optionally change the default anonymous UID and GID for the share. If the permission of a host is read-only or read-write (without allowing root access), and the UID of the client is 0 (which is typically the UID of the root account), then the UID is mapped to the anonymous UID on the NAS server. By default, the values of the anonymous UID and anonymous GID are 4294967294, which is typically associated with the `nobody` user.
5. On the **Access** page, optionally specify the name of the hosts that can access the share, along with their access privileges. In the **Default Access** field, select the access setting you want all hosts to have for the share. In the **Customize access for the following hosts** section do either of the following:
- Change the access privileges for existing hosts.
 - Add new hosts and specify individual access privileges for those hosts.

Change NFS share properties

Procedure

1. Under **Storage**, select **File > NFS Shares**.
2. Select the relevant NFS share, and then click the **Edit** icon.
3. On the **General** tab, change the share description.
4. On the **Host Access** tab, configure or add host access to the NFS share.

CHAPTER 5

Performance metrics for NFS

- [View historical performance metrics](#) 44
- [View real-time performance metrics](#)..... 44
- [File System Client Bandwidth](#).....44
- [File System Client Response Time](#)..... 45
- [File System Client I/O Size](#)..... 45
- [File System Client IOPS](#)..... 45
- [System - Client File System Bandwidth](#)..... 45
- [System - Client File System Response Time](#)..... 46
- [System - Client File System I/O Size](#)..... 46
- [System - Client File System IOPS](#)..... 46
- [System - NFS Bandwidth](#).....46
- [System - NFS I/O Size](#).....47
- [System - NFS IOPS](#).....47
- [System - NFS Response Time](#)..... 47
- [File System Bandwidth](#)..... 47
- [File System I/O Size](#)..... 48
- [File System IOPS](#)..... 48
- [System - File System Bandwidth](#).....48
- [System - File System I/O Size](#)..... 48
- [System - File System IOPS](#)..... 49
- [Tenant Bandwidth](#).....49

View historical performance metrics

Procedure

1. Under **System**, select **Performance**.
2. Select the historical metrics dashboard for the system for which you created a performance metrics display.
3. For each system dashboard, you can define the time range of the values displayed for all the metric line charts on that dashboard. The default time range is **Last 1 hour**. Alternatively, select one of the other time range values.

The time range selections are enabled only if Unisphere has data spanning that time range.

4. For a custom time range, select **Custom** and choose the start and end dates and times of the charts displayed. Click **OK**.
5. To drill down into the data displayed in the line chart, you can breakdown the data displayed into individual lines that show the categories and contributors that provide data to the performance metric. Choose among the breakdown categories available for a particular metric.

Each contributor displays as a different color line in the chart and is identified in the legend. You can quickly remove and add each contributor by clicking on its name in the legend. Use the breakdown display to determine if one contributor is adding to the aggregated total more than another contributor as well as analyze how a contributor's activity increases or decreases at a particular time.

6. Hover over a data point in the chart to display the date, time, and measurement associated with that data point. Gaps in metric data collection are displayed as gaps in the line chart.
7. For object-level line charts, such as those for LUNS, file systems, drives, and so forth, you can select **Percentage View** to view the data points as percentage values instead of absolute values.

View real-time performance metrics

Procedure

1. Under **System**, select **Performance**.
2. Select the real-time metrics dashboard for the system for which you created a performance metrics display.
3. Hover over a data point in the chart to display the date, time, and measurement associated with that data point. Gaps in metric data collection are displayed as gaps in the line chart.
4. For object-level line charts, such as those for LUNS, file systems, drives, and so forth, you can select **Percentage View** to view the data points as percentage values instead of absolute values.

File System Client Bandwidth

Total amount of file system client I/O requests, in KB/s, for the selected file systems.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

File System Client Response Time

Average time spent completing file system client I/O requests, in microseconds, for the selected file systems.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

File System Client I/O Size

Average size of file system client I/O requests, in KB, for the selected file systems.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

File System Client IOPS

Total number of file system client I/O requests, in I/O per second, for the selected file systems.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

System - Client File System Bandwidth

Total amount of file system client I/O requests, in KB/s, across all file systems in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - Client File System Response Time

Average time spent completing file system client I/O requests, in microseconds, across file systems in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - Client File System I/O Size

Average size of file system client I/O requests, in KB, across all file systems in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - Client File System IOPS

Total number of file system client I/O requests, in I/O per second, across all file systems in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - NFS Bandwidth

Total amount of NFS I/O requests, in KB/s, across all ports in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - NFS I/O Size

Average size of NFS I/O requests, in KB, across all ports in the storage system. Calculated as a weighted average, which gives more weight to the SP with the highest number of NFS I/O requests.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - NFS IOPS

Total number of NFS I/O requests, in I/O per second, across all ports in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - NFS Response Time

Average time spent completing NFS I/O requests, in microseconds, across all file systems in the storage system. Calculated as a weighted average, which gives more weight to the LUNs with the highest number of I/O requests.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

File System Bandwidth

Total amount of internal I/O requests, in KB/s, for the selected file systems.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

File System I/O Size

Average size of internal I/O requests, in KB, for the selected file systems.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

File System IOPS

Total number of internal I/O requests, in I/O per second, for the selected file systems.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

System - File System Bandwidth

Total amount of internal I/O requests, in KB/s, across all file systems in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - File System I/O Size

Average size of internal I/O requests, in KB, across all file systems in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

System - File System IOPS

Total number of internal I/O requests, in I/O per second, across all file systems in the storage system.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following categories:

This category...	Groups or filters data by the...
Storage Processor	SPs that processed the network traffic.
Read/Write	Traffic types: read I/O and write I/O.

Tenant Bandwidth

Total amount of I/O requests, in KB/s, for the selected tenant.

Breakdown and filter categories

The aggregated data can be broken down or filtered by the following category:

This category...	Groups or filters data by the...
Read/Write	Traffic types: read I/O and write I/O.

PART 2

Advanced functionality

[Chapter 6, "Managing quotas"](#)

[Chapter 7, "Configure IP routes"](#)

[Chapter 8, "Configuring IP multi-tenancy"](#)

[Chapter 9, "Troubleshooting an NFS configuration"](#)

CHAPTER 6

Managing quotas

- [About file system quotas.....](#)54
- [Recommended approach for configuring quotas.....](#)55
- [Quota policies.....](#) 55
- [Enable or disable the enforcement of user quotas on a quota tree.....](#)56
- [Enable or disable the enforcement of user quotas on a file system.....](#) 56
- [Create a user quota on a file system.....](#) 57
- [Create a quota tree on a file system.....](#)57
- [Create a user quota on a quota tree.....](#)58
- [View file system storage space usage by user.....](#)58
- [View quota tree storage space usage.....](#)58
- [Change quota properties for a file system.....](#)58
- [Change properties for a quota tree.....](#) 59
- [Change the quota policy for a file system.....](#) 60

About file system quotas

You can track and limit drive space consumption by configuring quotas for file systems at the file system or directory level. You can enable or disable quotas at any time, but it is recommended that you enable or disable them during non-peak production hours to avoid impacting file system operations.

Note

You cannot create quotas for read-only file systems.

Quota configurations

The storage system supports three types of quota configurations:

Quota configuration	Description
User quota on a file system	Limits the amount of storage consumed by an individual user storing data on the file system.
Quota on a directory (called a quota tree once a quota is applied)	<p>Limits the total amount of storage consumed on the directory. You can use quota trees to:</p> <ul style="list-style-type: none"> Set storage limits on a project basis. For example, you can establish quota trees for a project directory that has multiple users sharing and creating files in it. Track directory usage by setting the tree quota's hard and soft limits to 0 (zero). <hr/> <p>Note</p> <p>If you change the limits for a quota tree, the changes take effect immediately, without disrupting file system operations.</p>
User quota on a quota tree	Limits the amount of storage consumed by an individual user storing data on the quota tree.

Soft and hard limits

A quota can have a soft limit, hard limit, or both.

- A soft limit is a preferred limit on storage usage. The system issues a warning when a soft limit is reached. You can set a grace period for a file system or a quota tree, which counts down time once the soft limit is met. If the grace period expires, users cannot write to the file system or quota tree until more space becomes available, even if the hard limit has not been met.

Note

If you update the grace period value, the new value affects only the quota or quotas which will exceed the soft limit after the update is performed. Any existing quotas which have been counting down using the older grace period value will not be affected.

- A hard limit is an absolute limit on storage usage.

If a hard limit is reached for a user quota on a file system or quota tree, the user will not be able to write data to the file system or tree until more space becomes available. If a hard limit is reached for a quota tree, no user will be able to write data to the tree until more space becomes available.

Recommended approach for configuring quotas

It is recommended that you configure quotas before the storage system becomes active in a production environment, and that you follow this basic procedure:

1. Create a file system.
2. Determine which quota policy best suits the file system's environment, and select that policy. The default policy is File Size, which calculates drive usage in terms of logical file sizes, and ignores the size of directories and symbolic links.
3. Enable the enforcement of user quotas at the file system level, and define default limits for those quotas. If default limits are not specified, the system sets no drive usage limits for users, unless explicit user limits are defined for each individual user. Set default quotas in an environment where you want the same set of limits applied to many users.
4. Specify the grace period for which users of the file system can remain over the soft limit before it becomes the hard limit.

Note

If you update the grace period value, the new value affects only the quota or quotas which will exceed the soft limit after the update is performed. Any existing quotas which have been counting down using the older grace period value will not be affected.

-
5. Define explicit quotas for individual users at the file-system level, if the environment requires this type of usage-control granularity. The explicit quotas you define supersede the default quota definitions.
 6. Create quota trees for each directory or subdirectory for which you want to have quotas.
 7. For each quota tree, optionally change the default limits for users at the quota tree level. These limits are inherited from file system settings when a quota trees is created. If default limits are not set, the quotas feature sets no drive usage limits for quota tree users, unless explicit user limits are defined for each individual user. Set default limits in an environment where you want the same set of limits applied to many users.
 8. For each quota tree, define explicit quotas for users if the environment requires this type of individual-usage-control granularity.

Quota policies

Before enabling and defining quotas, ensure that the file system is configured to use the quota policy that best suits the client environment:

- **File Size policy (default):** Calculates drive usage in terms of logical file sizes, and ignores the size of directories and symbolic links. Use this policy where file sizes are critical to quotas, such as where user usage is based on the size of the files created, and exceeding the size limit is unacceptable.
- **Blocks policy:** Calculates drive usage in terms of file system blocks (8 KB units), and includes drive usage by directories and symbolic links in the calculations. With

this policy, any operation resulting in allocating or removing blocks, such as creating, expanding, or deleting a directory; writing or deleting files; or creating or deleting symbolic links changes block usage. Block usage depends solely on the number of bytes added to or removed from the file.

Note

When using the Blocks policy, a user can create a sparse file whose size is larger than the file size, but that uses fewer blocks on the drive.

The policy and grace period to use depend on which behavior (of the two described above) is preferred or the number of each type of client in your environment. If the grace period is set to 0, warnings will be generated when soft quotas are reached, but neither client will get quota exceeded errors until the hard limit is exceeded. If the use of default soft quotas is required, set the specific grace periods you desire, or keep the default grace period of one week.

Note

If you update a grace period value, the new value affects only the quota or quotas which will exceed the soft limit after the update is performed. Any existing quotas which have been counting down using the older grace period value will not be affected.

Enable or disable the enforcement of user quotas on a quota tree

Enabling or disabling the enforcement of user quotas on a quota tree impacts system performance, but does not disrupt file system operations. It is recommended that you perform these operations only during non-peak production hours. Once user quota enforcement is enabled, you can change quota settings without impacting performance.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then select the **Edit** icon.
3. On the **Quota** tab, select the **Quota Tree** sub-tab.
4. Do either of the following:
 - To enforce user quotas, locate the quota tree, and select the **No** link in the **Enforce User Quotas** column. Then select **Enforce User Quotas**.
 - To disable the enforcement of user quotas, locate the quota tree, and select the **Yes** link in the **Enforce User Quotas** column. Then clear **Enforce User Quotas**.

Enable or disable the enforcement of user quotas on a file system

Enabling or disabling the enforcement of user quotas on a file system impacts system performance, but does not disrupt file system operations. It is recommended that you perform these operations only during non-peak production hours. Once user quota

enforcement is enabled, you can change quota settings without impacting performance.

Note

When you enable user quotas, you can also set default user quota limits and a default grace period. Explicit user quotas will override these defaults.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then select the **Edit** icon.
3. On the **Quota** tab, select the **File System** sub-tab.
4. Select **Manage Quota Settings**.
5. Select or clear **Enforce User Quotas**.
6. If you are enabling user quotas, optionally do the following:
 - Change the quota policy for the file system.
 - Change the default quota limits and grace period. These limits apply to all file system users who do not have explicit user quotas defined. A value of 0 indicates no limit.

Create a user quota on a file system

Create a user quota on a file system to limit or track the amount of storage space that individual users consume on that file system. When you create or modify user quotas, you have the option to use default hard and soft limits that are set at the file-system level.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then select the **Edit** icon.
3. On the **Quota** tab, select the **File System** sub-tab.
4. Select the **Add** icon.
5. In the **Create User Quota** wizard, select the **Add** icon, and then provide the requested information. To track space consumption without setting limits, set **Soft Limit** and **Hard Limit** to 0, which indicates no limit.

Create a quota tree on a file system

Create a quota tree at the directory level of a file system to limit or track the total storage space consumed for that directory.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then select the **Edit** icon.
3. On the **Quota** tab, select the **Quota Tree** sub-tab.
4. Select the **Add** icon.

5. Follow the steps in the wizard. To track space consumption without setting limits, set the **Soft Limit** and **Hard Limit** fields to 0, which indicates no limit.

Create a user quota on a quota tree

Create a user quota on a quota tree to limit or track the amount of storage space that individual users consume on that tree. When you create user quotas, you have the option to use the default hard and soft limits that are set at the quota-tree level.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then select the **Edit** icon.
3. On the **Quota** tab, select the **Quota Tree** sub-tab.
4. Select the quota tree, and then select the **Edit** icon.
5. On the **User Quotas** tab, be sure that **Enforce User Quotas** is selected, and provide the limits information. To track space consumption without setting limits, set the **Soft Limit** and **Hard Limit** fields to 0, which indicates no limit.

View file system storage space usage by user

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then click the **Edit** icon.
3. Select the **Quota** tab to view the User Quota Report.

View quota tree storage space usage

You can view total quota tree storage space usage or quota tree space usage by user.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then click the **Edit** icon.
3. On the **Quota** tab, select the **Quota Tree** sub-tab.
The system displays the total storage space usage by quota tree.
4. To view quota tree storage space usage by user, select the quota tree, select the **Edit** icon, and then select the **User Quotas** tab.

Change quota properties for a file system

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the relevant file system, and then select the **Edit** icon.
3. On the **Quota** tab, select the **File System** sub-tab.
4. Change the limit settings for a user quota by selecting the quota and then selecting the **Edit** icon.

5. Select **Manage Quota Settings**, and do any of the following:
 - Change the quota policy for the file system.
 - Enforce user quotas on the file system.
 - Change the default soft limit, hard limit, and grace period for new user quotas on the file system. You can change these values for individual user quotas when you create them or when you modify their properties.

Note

If you update the grace period value, the new value affects only the quota or quotas which will exceed the soft limit after the update is performed. Any existing quotas which have been counting down using the older grace period value will not be affected.

Change properties for a quota tree

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the relevant file system, and then select the **Edit** icon.
3. On the **Quota** tab, select the **Quota Tree** sub-tab.
4. Select the relevant quota tree, and then select the **Edit** icon.
5. On the **General** tab, do any of the following:
 - Select **Use Default Limits** to keep the file system's default limits and grace period on the quota tree.
 - Clear **Use Default Limits** to override the file system's default limits and grace period.
 - Change the hard limit, soft limit, and grace period settings for the quota tree.

Note

If you update the grace period value, the new value affects only the quota or quotas which will exceed the soft limit after the update is performed. Any existing quotas which have been counting down using the older grace period value will not be affected.

6. On the **User Quotas** tab, do any of the following:
 - Select or clear **Enforce User Quotas** to enable or disable the enforcement of user quotas on the quota tree.
These actions impact system performance, but do not disrupt file system operations. It is recommended that you perform these operations only during non-peak production hours. You can change other quota settings without impacting performance.
 - If you enable the enforcement of user quotas on the quota tree, you can specify the soft and hard limits for those quotas. (You can override these values when you create individual quotas.)
 - Create a new user quota on the quota tree.

- Edit properties for existing user quotas.

Change the quota policy for a file system

Changing the quota policy for a file system can impact system performance, because it causes a system rescan. Therefore, it is recommended that you perform this action during off-peak hours.

Procedure

1. Under **Storage**, select **File > File Systems**.
2. Select the file system, and then select the **Edit** icon.
3. On the **Quota** tab, select **Manage Quota Settings**.
4. Change the quota policy, as desired.

CHAPTER 7

Configure IP routes

- [About NAS server routing](#) 62
- [NAS server interfaces](#) 64
- [Preferred interfaces for NAS servers](#) 64
- [IP Packet reflect functionality for NAS server interfaces](#) 65
- [Manage NAS server network interfaces and default routes](#) 65
- [Manage NAS server routes for responding to client requests](#) 66
- [Manage NAS server routes for external service requests](#) 66
- [Enable or disable IP packet reflect for a NAS server](#) 67
- [Verify NAS server routes](#) 67

About NAS server routing

You configure the IP interfaces and routing settings independently for each NAS server.

Configuring routes for responding to client requests

There are two ways to configure the routes for responding to client requests:

- Configure routing with IP packet reflect enabled.
- Configure routing with IP packet reflect disabled.

Every outbound packet sent in response to a client request always exits through the same interface that the inbound request used. This does not depend on IP packet reflect settings.

When IP packet reflect is enabled, you do not have to configure routing to clients that connect to the storage system, because the reply packets are sent back to the host or router where the packets came from. IP packet reflect is disabled by default.

Note

Requests that originate from the Unity system cannot leverage IP packet reflect, so you may still need to configure routing for external services, such as DNS and LDAP, when IP packet reflect is enabled.

When IP packet reflect is disabled, each NAS server interface uses static routing for directing packets to their destinations. To configure routes for responding to client requests, use the per-interface routing table, which is located by selecting **Show external routes for interfaces** on the **Network** tab of the NAS server properties page. You can add, modify, and delete routes in this table. Each route in the routing table directs a packet from the NAS server interface to which the route is linked.

Note

With static routing, the system does not check the link status or router availability. IP packet reflect, however, provides a return response that uses the request path of the client, without regard to the servers default or statically configured routes. If there is a router failure, replacement, or IP change, IP packet reflect supports the correct routing without interrupting the client connection.

Configuring routes to external services

In most cases, the NAS server interfaces are configured with a default gateway, which is used to route requests from a NAS server's interface to external services. You can add or view the default gateway for each NAS server interface by accessing the **External Services Access Routes** table. To access this table, select the **Routes to External Services** sub-tab on the **Network** tab of the NAS server properties page.

You can add or view default gateways by accessing the **Manage Routes** page, which displays all routes configured for the storage system in one place. To access this page, select the **Settings** icon, and then select **Access > Routing**.

You can add additional routes to these tables, as you would to any standard routing table, and you can modify or delete existing routes. When you make changes to routes in one table, the changes are reflected in the other table.

In a complex environment, you may need to configure granular routes to external services. To access a server from a specific interface through a specific gateway, add

a route with the following information following to the **External Services Access Routes** table:

```
From: <interface_ip>
Type: host
Gateway: <gateway_ip>
Destination: <external_server_ip>
Netmask/Prefix Length: 255.255.255.0
```

For example, to configure resilient DNS access, the standard recommendation is to configure the NAS server with three DNS servers, with each being accessed by a different physical or virtual connection. To do this:

- Add three DNS server IP addresses to the NAS server DNS configuration.
- Configure three NAS server interfaces, with each on a different physical port and/or VLAN.
- Add three routes as shown above, with each using a different NAS server interface IP and a different DNS server IP.

To access a server located on a different subnet, add a route like the following with the following information to the **External Services Access Routes** table.

```
From: <interface_ip>
Type: net
Gateway: empty
Destination: < subnet number>
Netmask/Prefix Length: <length>
```

NAS server routing tables

The per-interface routing table specifies routes from NAS server interfaces to client hosts. The system logic for picking the route of the per-interface table follows these rules:

- The routes are chosen from the NAS server's interfaces.
- The chosen interface must be active.
- If there are multiple routes to the same destination, the route specified by the preferred interface is chosen.
- If there are multiple routes to the same destination and there is no preferred interface, the most specific route takes precedence over the other routes. The order of precedence is host, net, default, with host being the most specific

The **External Services Access Routes** table is dynamically created by merging the per-interface routing tables with preferred interface information. The system chooses the best possible routing configuration when NAS server interfaces are added, modified, or deleted, either manually or through replication changes. The system logic for picking the route of the **External Services Access Routes** table follows these rules:

- The routes are chosen from the NAS server's interfaces.
- If there are multiple routes to the same destination, the route specified by the preferred interface is chosen.
- If there are multiple routes to the same destination and there is no preferred interface, the most specific route takes precedence over the other routes. The order of precedence is host, net, default, with host being the most specific

For both routing tables, the system logic also contains algorithms for handling more complicated configurations.

NAS server interfaces

When you modify an IP interface for a NAS server, you can specify whether it:

- Is a production or backup interface.
- Is a preferred interface, which is used for outgoing communication with non-locally connected hosts.

Preferred interfaces for NAS servers

If you have multiple interfaces configured for a NAS server, the system will automatically select the interface that the default route uses for outgoing communication to external services. To change which interface is selected, you can specify preferred interface settings.

The NAS server uses preferred interfaces in the following circumstances:

- The application does not specify the source interface.
- The destination is on a remote subnet.

Note

Locally connected hosts, which are attached to the same subnets as the NAS server interfaces, are accessed by using corresponding interfaces directly, and not through the preferred interface gateways.

You can select one preferred interface for each of the following interface types:

- IPv4 interface of type Production
- IPv6 interface of type Production
- IPv4 interface of type Backup & DR Testing
- IPv6 interface of type Backup & DR Testing

When the **Preferred Interface** field is set to **Auto** (the default), the system selects the preferred interface automatically, based on how many routes the interface has and how wide the destination range is of its routes. For most user environments using **Auto** provides an optimal selection of the preferred interface.

When a NAS server initiates outbound traffic to an external service, it compiles a list of all the available network interfaces on the proper subnet and performs one of the following actions if a preferred interface of the appropriate type (IPv4 or IPv6) is in the compiled list:

- If the preferred production interface is active, the system uses the preferred production interface.
- If the preferred production interface is not active, and there is a preferred active backup interface, the system uses the preferred backup interface.
- If the preferred production interface is not active (as in the case of a NAS server failover), and there is no preferred backup interface, the system does nothing.

If a preferred interface is not in the compiled list, the underlying operating environment platform chooses the network interface.

IP Packet reflect functionality for NAS server interfaces

IP packet reflect functionality for NAS servers ensures that outbound (reply) packets always exit through the next hop gateway through which inbound (request) packets entered. Because the majority of network traffic on a NAS server (including all file system I/O) is client-initiated, the NAS server can use IP packet reflect to reply to client requests. IP packet reflect is disabled by default.

Note

Interface selection is not affected by IP packet reflect settings.

IP packet reflect provides the following advantages:

- With IP packet reflect, there is no need to determine the route for sending the reply packets.
- Improves network security. Because reply packets always go out the same next hop gateway as the request packets, request packets cannot be used to indirectly flood other LANs. In cases where two network devices exist, one connected to the Internet and the other connected to the intranet, replies to Internet requests do not appear on the intranet.
- Supports multiple subnets, with each on a different NIC. With this configuration, each subnet uses a router, and the router port for each subnet filters incoming packets, so only packets from that subnet are forwarded. Replies, therefore, must be sent through the same next hop gateway as the incoming requests. IP packet reflect satisfies this requirement.
- Helps clients that have a single IP address and multiple MAC addresses. Although unusual, this configuration creates a problem for the server if IP packet reflect is not enabled. For each IP address, the NAS server keeps only one associated MAC address in the Address Resolution Protocol (ARP) table. With IP packet reflect enabled, this problem is resolved, because the server does not need to look up the MAC address from the ARP database for the reply. Instead, the server uses the MAC address of the request to send the reply.

Manage NAS server network interfaces and default routes

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the NAS server to modify, and select the **Edit** icon.
3. Select the **Network** tab.
4. Make the desired changes as follows:

Task	Description
Add a network interface and default route	<ol style="list-style-type: none"> a. In the Network Interfaces field, select the Add icon, and then select the type of IP interface to add. b. Select the port and enter the IP address for the new interface. c. Optionally enter a gateway to use for the default route.

Task	Description
	d. If the switch port supports VLAN tagging, optionally specify a VLAN ID (between 0 and 4095) for the VLAN with which the NAS server is associated. If the NAS server is associated with a tenant, you must select a VLAN ID.
Modify a network interface	<p>a. In the Network Interfaces field, select the network interface to modify, and then select the Edit icon.</p> <p>b. Modify the desired values.</p>
Specify or change the preferred network interfaces	<p>a. Select Change Preferred Interface.</p> <p>b. Select the appropriate preferred interfaces or select Auto.</p>
Remove a network interface	<p>Select the network interface you wish to remove from the NAS Server configuration, and click the Delete icon.</p> <hr/> <p>Note</p> <p>If you delete a preferred interface, the system will select a new preferred interface.</p> <hr/>

Manage NAS server routes for responding to client requests

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the NAS server, and then select the **Edit** icon.
3. Select the **Network** tab, and then select the interfaces for which you are configuring routes.
4. Select **Show external routes for interfaces**, near the bottom of the screen.
5. To add a route, select the **Add** icon in the per-interface routing table, and then specify the relevant information.
6. To change a route, follow these steps.
 - a. Select the interface in the network interfaces table.
 - b. Select the route and select the **Edit** icon in the per-interface routing table.
 - c. Specify the relevant information.

Manage NAS server routes for external service requests

Routes for external service requests are routes that the system uses to request external services, such as LDAP or DNS.

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the NAS server, and then select the **Edit** icon.
3. Select the **Network** tab.
4. Select **Routes to External Services**.

5. To add a route, select the **Add** icon, and then specify the relevant information.
6. To change a route, select the route, select the **Edit** icon, and then specify the relevant information.
7. To hide default and local subnet routes from view, select **More Actions > Hide default and local subnet routes**.

Enable or disable IP packet reflect for a NAS server

Before you begin

You can enable or disable IP packet reflect for each NAS server. IP packet reflect is disabled for all NAS servers by default.

Before you disable IP packet reflect, make sure that the hosts are reachable through a default, network, or host route. Otherwise, some hosts may become unavailable when IP packet reflect is disabled.

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the NAS server to modify, and select the **Edit** icon.
3. Select the **Network** tab.
4. In the **Packet Reflect** field, select the **Edit** icon, and then select **Enabled** or **Disabled**.

Verify NAS server routes

You can verify NAS server routes using the Ping and Trace operations. You can verify routes from all system interfaces, except the management interface.

Procedure

1. Under **Storage**, select **File > NAS Servers**.
2. Select the NAS server, and then select the **Edit** icon.
3. Select the **Network** tab.
4. To verify routes from a specific interface:
 - a. From the **Interfaces & Routes** sub-tab, select the interface, and then select **Ping/Trace**.
 - b. Fill in the requested information, and select **Ping** or **Trace**.
5. To verify routes from any interface:
 - a. Select the **Routes to External Services** sub-tab, and then select **Ping/Trace**.
 - b. Fill in the requested information, and select **Ping** or **Trace**.

Configure IP routes

CHAPTER 8

Configuring IP multi-tenancy

- [About IP multi-tenancy](#) 70
- [Configuring IP multi-tenancy](#) 70
- [Add a tenant](#) 71
- [Change tenant properties](#) 71
- [Configure file replication for a tenant](#) 72

About IP multi-tenancy

IP multi-tenancy provides the ability to assign isolated, file-based storage partitions to the NAS servers on a storage processor. Tenants are used to enable the cost-effective management of available resources, while at the same time ensuring that tenant visibility and management is restricted to assigned resources only.

With IP multi-tenancy, each tenant can have its own:

- IP addresses and port numbers.
- VLAN domain.
- Routing table.
- IP firewall.
- DNS server or other administrative servers to allow the tenant to have its own authentication and security validation.

IP multi-tenancy is implemented by adding a tenant to the storage system, associating a set of VLANs with the tenant, and then creating one NAS server for each of the tenant's VLANs, as needed. It is recommended that you create a separate pool for the tenant and that you associate that pool with all of the tenant's NAS servers.

Note the following about the IP multi-tenancy feature:

- There is a one-to-many relationship between tenants and NAS servers. A tenant can be associated with multiple NAS servers, but a NAS server can be associated with only one tenant.
- You can associate a NAS server with a tenant when you create the NAS server. Once you create a NAS server that is associated with a tenant, you cannot change this association. (You cannot associate this NAS server with any other tenant or remove the association with this tenant.)
- During replication, data for a tenant is transferred over the service provider's network rather than the tenant's network.
- Because multiple tenants can share the same storage system, a spike in traffic for one tenant can negatively impact the response time for other tenants.

Configuring IP multi-tenancy

To configure IP multi-tenancy, follow this process:

1. Create a storage pool for each tenant (recommended).
2. Add the tenants to the system. When you add tenants, you assign each one a non-overlapping set of VLANs.
3. Create a NAS server for each tenant. When you create a NAS server, select the tenant to associate with the NAS server, and select the tenant's pool, which will be used to store the NAS server's metadata. You can add network interface information for the tenant now or later on.

Note

In a network interface, each subnet must be unique for a given VLAN. Using the same subnet for different VLANs can cause connectivity issues.

4. Create the file systems and shares for each tenant.

5. Configure hosts access for the tenant's NFS shares.

Example

The following table shows the Unity components used for tenants T1 and T2. In this example, each tenant has two VLANs and separate NAS servers for the Engineering (eng) and Human Resources (hr) departments. Each NAS server has one file system and one share.

Table 5 Unity components for tenant T1

Pool	VLANs	NAS servers	File systems	Shares
T1_pool	900	T1_nfs_eng	T1_nfs_eng_fs	T1_nfs_eng_sh
	901	T1_nfs_hr	T1_nfs_hr_fs	T1_nfs_hr_sh

Table 6 Unity components for tenant T2

Pool	VLANs	NAS servers	File systems	Shares
T2_pool	902	T2_nfs_eng	T2_nfs_eng_fs	T2_nfs_eng_sh
	903	T2_nfs_hr	T2_nfs_hr_fs	T2_nfs_hr_sh

Add a tenant

Before you begin

Obtain the VLAN IDs to associate with the tenant.

Procedure

1. Under **Storage**, select **File > Tenants**.
2. Select the **Add** icon.
3. Specify the information on the **Add Tenant** window. If this is the first creation of a tenant in your environment, have the system automatically generate a UUID value for this tenant. Otherwise, for existing tenants in your environment that have a system generated UUID value, enter that UUID value manually.

Change tenant properties

Procedure

1. Under **Storage**, select **File > Tenants**.
2. Select the **Edit** icon.
3. Change the tenant name, and add or remove associated VLANs. You can add a VLAN ID to a tenant if:
 - The VLAN ID is not associated with an existing tenant.
 - No network interfaces use the VLAN ID.

Configure file replication for a tenant

In a multi-tenancy environment, you can replicate the NAS servers, routes, and file systems for a specific tenant.

For general information about replication, see the Unity online help and *Configuring Replication*, which is available from the [UnityOE Features Info Hub](#).

Procedure

1. Create a pool for the tenant on the destination system.
2. Add the tenant to the destination system. When you add the tenant, use the same UUID and VLANs as the tenant on the source system.
3. If you are configuring remote replication, perform the following steps to set up the remote connection. Once you set this up, the same connection can be used again for subsequent replication sessions between the same systems.
 - a. Configure a mobility interface on the source and destination systems. The IP addresses of both systems should be on the same subnet.
 - b. Configure a replication connection on the source system using the **Asynchronous** connection mode.
4. On the NAS server properties page, create a replication session for the NAS server associated with the file storage. When you configure this session, specify the pool you created in Step 1.

Storage resources included in a NAS server automatically get replicated when a replication session is first configured for the NAS server. The replication session for the storage resources will inherit the same attributes as the associated replication session of the associated NAS server. For the storage resources you do not want participating in replication, you can choose to remove the associated replication sessions manually.

5. To configure automatic synchronization of the NAS server and all of its files, select **Sync** on the **Replication** tab of the source NAS server.
6. To replicate the NAS server and a specific file system, access the properties page for the source file system, and select **Sync** on the **Replication** tab.

CHAPTER 9

Troubleshooting an NFS configuration

- [Service commands for troubleshooting NFS issues in Unity](#)..... 74

Service commands for troubleshooting NFS issues in Unity

The following service commands are useful for troubleshooting NFS issues in Unity. For detailed information about service commands, see the *Service Commands Technical Notes*, which is available from the [Unity All-Flash & Hybrid Info Hub](#).

Use case	Service command
Perform NAS server advanced management. This includes displaying and customizing the parameters of various NAS components, performing database maintenance, and performing network troubleshooting.	<code>svc_nas</code>
Display the settings and server connection status for the Common Event Publishing Agent for a specified NAS server.	<code>svc_event_publishing</code>
Display or reset the counters for NDMP and PAX backup statistics.	<code>svc_pax</code>
View information about locks currently held for provisioned Unity storage.	<code>svc_lockd</code>
Dump the VHDX metadata (Hyper-V virtual disk files) to diagnose issues with VHDX files.	<code>svc_vhdx</code>