



Hitachi NAS Platform

Server and Cluster Administration Guide

Release 12.3

© 2011-2015 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd.

Hitachi, Ltd., reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. This document contains the most current information available at the time of publication. When new or revised information becomes available, this entire document will be updated and distributed to all registered users.

Some of the features described in this document might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Data Systems Corporation at <https://portal.hds.com>.

Notice: Hitachi, Ltd., products and services can be ordered only under the terms and conditions of the applicable Hitachi Data Systems Corporation agreements. The use of Hitachi, Ltd., products is governed by the terms of your agreements with Hitachi Data Systems Corporation.

Hitachi Data Systems products and services can be ordered only under the terms and conditions of Hitachi Data Systems' applicable agreements. The use of Hitachi Data Systems products is governed by the terms of your agreements with Hitachi Data Systems.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

Archivas, Dynamic Provisioning, Essential NAS Platform, HiCommand, Hi-Track, ShadowImage, Tagmaserve, Tagmasoft, Tagmasolve, Tagmastore, TrueCopy, Universal Star Network, and Universal Storage Platform are registered trademarks of Hitachi Data Systems Corporation.

AIX, AS/400, DB2, Domino, DS8000, Enterprise Storage Server, ESCON, FICON, FlashCopy, IBM, Lotus, OS/390, RS6000, S/390, System z9, System z10, Tivoli, VM/ESA, z/OS, z9, zSeries, z/VM, z/VSE are registered trademarks and DS6000, MVS, and z10 are trademarks of International Business Machines Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). Some parts of ADC use open source code from Network Appliance, Inc. and Traakan, Inc.

Part of the software embedded in this product is gSOAP software. Portions created by gSOAP are copyright 2001-2009 Robert A. Van Engelen, Genivia Inc. All rights reserved. The software in this product was in part provided by Genivia Inc. and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the author be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

The product described in this guide may be protected by one or more U.S. patents, foreign patents, or pending applications.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Contact the Hitachi Data Systems Legal Department for any export compliance questions.



Contents

Preface.....	10
Contacting Hitachi Data Systems.....	10
Related Documentation.....	10
1 Configuring the System Management Unit (SMU).....	14
Using the SMU Setup Wizard.....	15
Configuring SMU security.....	18
Disabling protocols and cipher suites.....	20
Configuring an SMTP relay for the SMU.....	21
Displaying the SMU software version.....	23
Upgrading SMU software and server firmware.....	23
Selecting SMU-managed servers.....	23
Changing the IP address of a managed server.....	26
Using the SMU as an NTP server.....	28
2 Configuring the storage server.....	30
Cloning server settings.....	31
Using the Server Setup Wizard.....	33
Configuring server management access.....	34
Setting the server password.....	34
Configuring SSC access.....	35
Configuring SNMP access.....	37
Configuring SNMPv3 access.....	39
Configuring server identification.....	40
Configuring server date and time.....	42
Storage server and NTP server interaction.....	42
Configuring storage server date and time.....	42
Managing license keys.....	44
Adding a license key.....	46
License types.....	47
Displaying storage server version information.....	50

Upgrading SMU software and server firmware.....	52
3 Clustering.....	54
Clusters and server farms.....	56
Clusters.....	56
Nway clustering.....	56
Maximum number of nodes supported.....	57
Quorum device (QD) in a cluster configuration.....	58
Cluster topology.....	59
Enhanced cluster quorum device.....	59
Server farms.....	59
Clusters versus server farms.....	61
Using clusters.....	61
Cluster name space (CNS).....	62
EVS name spaces.....	63
About cluster licensing.....	63
Configuring a new cluster.....	65
Configuring the first cluster node.....	65
Joining an existing cluster using Web Manager.....	66
Configuring the cluster.....	67
Displaying cluster node details.....	68
Quorum device management (external SMUs only).....	72
Using cluster name space (CNS).....	73
CNS usage considerations.....	73
Displaying the cluster name space tree.....	74
Displaying the EVS name space tree.....	75
Managing links and subdirectories in the EVS name space.....	75
Creating a cluster name space tree.....	75
Creating a CNS root directory.....	76
Creating CNS subdirectories.....	76
Creating a file system link.....	76
Changing cluster name space properties.....	78
Deleting a cluster name space.....	78
Renaming a CNS subdirectory.....	78
Moving a CNS directory.....	78
Deleting a CNS directory.....	79
Modifying a file system link.....	79
Deleting a file system link.....	79
Configuring read caching.....	80
Configuring file caching options.....	81
Reviewing read cache statistics.....	82
Displaying read cache statistics.....	84
Deleting a read cache.....	85
Read cache considerations.....	85
4 Using virtual servers (EVSs).....	88
Secure virtual servers.....	89
Secure EVS considerations.....	89
About security contexts.....	91
Security context contents.....	92

Securing an EVS.....	92
Removing an individual security context from a secure EVS.....	95
EVS name spaces.....	96
Creating an EVS.....	97
Assigning a file system to an EVS.....	97
Virtual server (EVS) management.....	98
Displaying EVS details.....	101
Migrating an EVS within a cluster.....	103
Migrating an EVS within a server farm.....	105
Cloning server settings.....	105
Migrating an EVS within a server farm.....	106
5 Status and monitoring.....	108
Storage system status.....	110
Configuring devices on the System Monitor.....	110
Checking the system status.....	112
Using the server status console.....	113
Checking the status of a server unit.....	116
Checking UPS status.....	120
Checking SMU status.....	120
Monitoring multiple servers.....	122
Monitoring storage subsystems with Hitachi Device Manager.....	123
Managing HDvM server connections.....	123
Connecting the SMU to an HDvM server.....	124
Changing HDvM server connection details.....	125
Removing HDvM server connections.....	126
6 Performance graphs.....	128
Available performance graphs.....	129
Controlling the performance graph display.....	130
Displaying a custom date range.....	130
Displaying the Performance Graphs page.....	132
Displaying Node Ops/Sec.....	134
Displaying Ethernet Throughput.....	136
Displaying System Load.....	138
Displaying Disk Latency.....	140
Displaying Fibre Channel Throughput.....	142
Displaying Cache and Heap Usage.....	144
Displaying NVRAM Waited Allocs.....	146
Displaying Running Bossock Fibers.....	148
Displaying File System Ops/Sec.....	150
Displaying File System Capacity.....	152
Displaying Storage Pool Capacity.....	154
Downloading performance data.....	156
Storage server statistics.....	157
Network statistics.....	157
Ethernet statistics.....	157
Displaying Ethernet Statistics.....	157
Displaying aggregated ports or per-port Ethernet statistics.....	159
TCP/IP Statistics.....	160

Displaying TCP/IP statistics.....	161
Displaying aggregated ports or per-port TCP/IP statistics.....	162
Displaying TCP/IP detailed statistics.....	164
Fibre Channel statistics.....	165
Displaying Fibre Channel statistics.....	166
Displaying per port Fibre Channel statistics.....	168
File and block protocol statistics.....	169
Displaying NFS statistics.....	169
Displaying CIFS statistics.....	172
Displaying FTP statistics.....	182
Displaying iSCSI statistics.....	184
Data access and performance statistics.....	187
Server and file system load statistics.....	187
Displaying operations per second (ops/sec) statistics.....	188
Displaying file system NVRAM statistics.....	189
Management statistics.....	189
Displaying access management statistics.....	191
Displaying SNMP management statistics.....	192
Displaying HTTPS management statistics.....	195
Displaying VSS management statistics.....	197
Displaying virus scanning statistics.....	199
Event logging and notification.....	201
Using the event log.....	202
Displaying and filtering the event log.....	202
Configuring event notifications.....	204
Using SNMP and syslog.....	215
Clearing logs with Windows Event Viewer.....	221
File system auditing.....	221
About file system audit logs.....	222
Controlling file system auditing.....	223
Creating a file system audit policy.....	223
Enabling auditing for a file system.....	226
Modifying a file system audit policy.....	228
Enabling or disabling auditing for a file system.....	230
Deleting a file system audit policy.....	231
Displaying file system audit logs.....	231
FTP auditing.....	232
Displaying FTP Audit Logs page.....	232
Enabling or disabling FTP audit logging for an EVS.....	233
Configuring FTP audit logging.....	234
Displaying FTP audit logs.....	236
Monitoring Fibre Channel switches.....	236
Displaying Fibre Channel switch connectivity status.....	237
Using System Monitor to display switch connectivity status.....	237
Using Web Manager to display switch connectivity status.....	238
Adding FC switches.....	238
Displaying or changing details for an FC switch.....	240
Optimizing performance with Performance Accelerator.....	241
Determining if Performance Accelerator will increase system performance.....	242
Installing Performance Accelerator.....	242
Uninstalling Performance Accelerator.....	242
Troubleshooting Performance Accelerator.....	243

7	Providing an SSL certificate to the external SMU.....	246
	Generating a custom private key and SSL certificate.....	247
	Generating a certificate signing request (CSR).....	248
	Installing certificates.....	248
	Recreating the default SMU certificate.....	249
	Accepting self-signing certificates.....	250
8	Providing an SSL certificate to the embedded SMU	254
	Configuring cipher suites.....	255
	Configuring the SSL/TLS version.....	256
	Obtaining and importing a CA-signed certificate.....	257
A	Using HNAS multi-tenancy.....	260
	Using HNAS multi-tenancy.....	261
	Understanding multi-tenancy.....	261
	Understanding HNAS multi-tenancy benefits.....	262
	How multi-tenancy mode differs from stand-alone mode.....	263
	How multi-tenancy differs from per-EVS security.....	263
	Multi-tenancy requirements.....	264
	Disabling HNAS multi-tenancy.....	265
	Managing multi-tenancy.....	265
	Multi-tenancy management interfaces.....	265
	Viewing HNAS multi-tenancy status	265
	Considerations for enabling HNAS multi-tenancy.....	266
	HNAS multi-tenancy limits.....	266
	Enabling HNAS multi-tenancy.....	266
	Managing multi-tenancy on the NAS server.....	267
	Managing multi-tenancy for an EVS.....	267
	Overlapping IP address support for HNAS multi-tenancy.....	268
	Understanding routing by EVS.....	269
	Configuring routes per EVS.....	269
	Understanding EVS crosstalk checking.....	271
	Multi-tenancy-aware protocols.....	271

Preface

In PDF format, this guide provides information about administering servers, clusters, and server farms. Includes information about licensing, name spaces, upgrading firmware, monitoring servers and clusters, the backing up and restoring configurations.

Contacting Hitachi Data Systems

2845 Lafayette Street
Santa Clara, California 95050-2627
U.S.A.
<https://portal.hds.com>
North America: 1-800-446-0744

Related Documentation

Release Notes provide the most up-to-date information about the system, including new feature summaries, upgrade instructions, and fixed and known defects.

Administration Guides

- *System Access Guide* (MK-92HNAS014)—In PDF format, this guide explains how to log in to the system, provides information about accessing the NAS server/cluster CLI and the SMU CLI, and provides information about the documentation, help, and search capabilities available in the system.
- *Server and Cluster Administration Guide* (MK-92HNAS010)—In PDF format, this guide provides information about administering servers, clusters, and server farms. Includes information about licensing, name spaces, upgrading firmware, monitoring servers and clusters, the backing up and restoring configurations.
- *Storage System User Administration Guide* (MK-92HNAS013)—In PDF format, this guide explains user management, including the different types of system administrator, their roles, and how to create and manage these users.
- *Network Administration Guide* (MK-92HNAS008)—In PDF format, this guide provides information about the server's network usage, and explains how to configure network interfaces, IP addressing, name and directory services.
- *File Services Administration Guide* (MK-92HNAS006)—In PDF format, this guide explains about file system formats, and provides information about

creating and managing file systems, and enabling and configuring file services (file service protocols).

- *Data Migrator Administration Guide* (MK-92HNAS005) —In PDF format, this guide provides information about the Data Migrator feature, including how to set up migration policies and schedules.
- *Storage Subsystem Administration Guide* (MK-92HNAS012)—In PDF format, this guide provides information about managing the supported storage subsystems (RAID arrays) attached to the server/cluster. Includes information about tiered storage, storage pools, system drives (SDs), SD groups, and other storage device related configuration and management features and functions.
- *Snapshot Administration Guide* (MK-92HNAS011)—In PDF format, this guide provides information about configuring the server to take and manage snapshots.
- *Replication and Disaster Recovery Administration Guide* (MK-92HNAS009) —In PDF format, this guide provides information about replicating data using file-based replication and object-based replication, provides information on setting up replication policies and schedules, and using replication features for disaster recovery purposes.
- *Antivirus Administration Guide* (MK-92HNAS004)—In PDF format, this guide describes the supported antivirus engines, provides information about how to enable them, and how to configure the system to use them.
- *Backup Administration Guide* (MK-92HNAS007)—In PDF format, this guide provides information about configuring the server to work with NDMP, and making and managing NDMP backups. Also includes information about Hitachi NAS Synchronous Image Backup.
- *Command Line Reference* Opens in a browser, and describes the commands used to administer the system.



Note: For a complete list of Hitachi NAS open source software copyrights and licenses, see the *System Access Guide*.

Hardware References

- *Hitachi NAS Platform 3080 and 3090 G1 Hardware Reference* (MK-92HNAS016)—Provides an overview of the second-generation server hardware, describes how to resolve any problems, and replace potentially faulty parts.
- *Hitachi NAS Platform 3080 and 3090 G2 Hardware Reference* (MK-92HNAS017)—Provides an overview of the second-generation server hardware, describes how to resolve any problems, and replace potentially faulty parts.
- *Hitachi NAS Platform Series 4000 Hardware Reference* (MK-92HNAS030) (MK-92HNAS030)—Provides an overview of the Hitachi NAS Platform Series 4000 server hardware, describes how to resolve any problems, and how to replace potentially faulty components.

- *Hitachi High-performance NAS Platform (MK-99BA012-13)*—Provides an overview of the NAS Platform 3100/NAS Platform 3200 server hardware, and describes how to resolve any problems, and replace potentially faulty parts.

Best Practices

- *Hitachi USP-V/VSP Best Practice Guide for HNAS Solutions (MK-92HNAS025)*—The HNAS practices outlined in this document describe how to configure the HNAS system to achieve the best results.
- *Hitachi Unified Storage VM Best Practices Guide for HNAS Solutions (MK-92HNAS026)*—The HNAS system is capable of heavily driving a storage array and disks. The HNAS practices outlined in this document describe how to configure the HNAS system to achieve the best results.
- *Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere (MK-92HNAS028)*—This document covers VMware best practices specific to HDS HNAS storage.
- *Hitachi NAS Platform Deduplication Best Practice (MK-92HNAS031)* —This document provides best practices and guidelines for using HNAS Deduplication.
- *Hitachi NAS Platform Best Practices for Tiered File Systems (MK-92HNAS038)* —This document describes the Hitachi NAS Platform feature that automatically and intelligently separates data and metadata onto different Tiers of storage called Tiered File Systems (TFS).
- *Hitachi NAS Platform Data Migrator to Cloud Best Practices Guide (MK-92HNAS045)*—Data Migrator to Cloud allows files hosted on the HNAS server to be transparently migrated to cloud storage, providing the benefits associated with both local and cloud storage.
- *Brocade VDX 6730 Switch Configuration for use in an HNAS Cluster Configuration Guide (MK-92HNAS046)*—This document describes how to configure a Brocade VDX 6730 switch for use as an ISL (inter-switch link) or an ICC (inter-cluster communication) switch.
- *Best Practices for Hitachi NAS Universal Migrator (MK-92HNAS047)*—The Hitachi NAS Universal Migrator (UM) feature provides customers with a convenient and minimally disruptive method to migrate from their existing NAS system to the Hitachi NAS Platform. The practices and recommendations outlined in this document describe how to best use this feature.
- *Hitachi NAS Platform Storage Pool and HDP Best Practices (MK-92HNAS048)*—This document details the best practices for configuring and using HNAS storage pools, related features, and Hitachi Dynamic Provisioning (HDP).
- *Hitachi Data Systems SU 12.x Network File System (NFS) Version 4 Feature Description (MK-92HNAS056)* —This document describes the features of Network File System (NFS) Version 4.

- *Hitachi NAS Platform Hitachi Dynamic Provisioning with HNAS v12.1* (MK-92HNAS057)—This document lists frequently asked questions regarding the use of Hitachi Dynamic Provisioning.
- *Hitachi Multi-tenancy Implementation and Best Practice Guide* (MK-92HNAS059)—This document details the best practices for configuring and using HNAS Multi-Tenancy and related features, and EVS security.

Configuring the System Management Unit (SMU)

The System Management Unit (SMU) manages the storage servers/clusters and controls data migration and replication policies and schedules. For example, you can:

- Secure the SMU, so that only certain predefined hosts can access the SMU for management purposes.
- Configure the SMU to act as an SMTP relay to the public network.



Note: For HTTP (Web Manager) access, the SMU ships with the default user name `admin` and the password `nasadmin`.

- ☐ [Using the SMU Setup Wizard](#)
- ☐ [Configuring SMU security](#)
- ☐ [Disabling protocols and cipher suites](#)
- ☐ [Configuring an SMTP relay for the SMU](#)
- ☐ [Displaying the SMU software version](#)
- ☐ [Upgrading SMU software and server firmware](#)
- ☐ [Selecting SMU-managed servers](#)
- ☐ [Changing the IP address of a managed server](#)
- ☐ [Using the SMU as an NTP server](#)

Using the SMU Setup Wizard

For an external SMU, basic SMU configuration is usually performed as a part of system initialization. The SMU setup wizard is used to complete the basic configuration of an SMU. Using the SMU Setup Wizard, you can change the administrator's password, set up name services for network operation, specify an SMTP server to relay email from the NAS server, and configure the date and time settings.

Procedure

1. Navigate to **Home > SMU Administration > SMU Setup Wizard** to change the administrator's password.

The default password is "nasadmin," and you should change it as soon as possible.

Field/Item	Description
Current Password	The current password for the currently logged in user.
New Password	The current password for the currently logged in user.
Confirm New Password	The current password for the currently logged in user.
next	Saves configuration changes, and proceeds to the next page of the wizard (DNS configuration).
cancel	Returns to the Home page without saving configuration changes.

2. When you have changed the password, click **next** to display the name services configuration page.

Field/Item	Description
DNS Servers	Enter the IP addresses of the DNS servers that will be applied to the SMU. <ul style="list-style-type: none">• A DNS server can be added by IPv4 or IPv6 address (never by host name).• If A DNS server is to be added by IPv6 address, IPv6 must be configured on the SMU.• A DNS server can resolve host names to IPv4 or IPv6 addresses, regardless whether it is connected to by IPv4 or IPv6.
DNS Domain	Displays the current domain.
Domain Search Order	Lists up to six domains that will be searched. To add a domain to the list, enter the DNS server name or IP address in the field, then click the add button (the down arrow). To remove a domain from the list, select the server to remove, then click the remove button (the X).

Field/Item	Description
	To change the domain search order, select the domain you want to move up or down in the list, then click the up or down arrow to change the domain's position in the list.
apply	Only displays on the Name Services page.
back	Returns to the password specification page.
next	Saves configuration changes, and proceeds to the next page of the wizard (SMTP configuration).
cancel	Returns to the Home page without saving configuration changes.

3. When you have specified the name services settings, click **next** to display the SMTP server configuration page.



Note: If you are accessing this page through the SMU Setup Wizard, the title of the page is **SMU Setup Wizard**, and there are three buttons at the bottom of the page (**back**, **next**, and **cancel**) to take you through the SMU setup.


Field/Item	Description
SMTP Server	Enter the SMTP Server on the public network. The SMU will then relay emails from the servers and other devices on the private network to the public network. An SMTP server can be specified by IPv4 or IPv6 address, or by a host name. If an IPv6 address is specified, the SMU will only be able to use the server for email forwarding if the SMU is configured with an IPv6 address. Additionally if the SMTP server is given by host name, and that host name resolves only to an IPv6 address, mail forwarding will only be possible if an IPv6 DNS server is provided.
apply	Saves configuration changes, and closes the page.

4. When you have specified the SMTP server, click **next** to display the date and time settings page.

Proper server operation requires time synchronization with a reliable time source. For example, Kerberos authentication (required when operating with Active Directory) depends on the current time. Clock 'drift' may also cause inaccurate reporting of file access and modification times, with unexpected results in data migrations. NTP provides the best and most reliable method for maintaining the server's time accuracy.



Note: If you are accessing this page through the **SMU Setup Wizard**, the title of the page is **SMU Setup Wizard**, and there are three buttons at the bottom of the page (**back**, **next**, and **cancel**) to take you through the SMU setup. If accessing through the **SMU Administration** page, the only button that appears is **apply**.

Field/Item	Description
Time	Current time (in 24-hour format).
Date	Current date (in YYYY-MM-DD format). Specify the current date using the calendar pop-up.
Time Zone	Time zone where the NAS server/cluster is located. Select the correct time zone from the drop-down list.
NTP Server IP/Name	<p>An NTP server can be specified by IPv4 or IPv6 address, or by a host name.</p> <p>When specified, the SMU gets the current date and time from this NTP server, and periodically checks with this NTP server to keep the SMU's clock accurate. If several NTP servers are specified, the SMU uses the first one in the list that it can contact.</p> <hr/> <p> Note: If an IPv6 address is specified, the SMU will only be able to synchronize if the SMU is configured with an IPv6 address. Additionally, if the NTP server is specified by host name, and that host name resolves to an IPv6 address, synchronization is only possible if an IPv6 DNS server is provided.</p> <hr/> <p>To add an NTP server, type the NTP server's DNS name or IP address in the field, then click the add button (the down arrow).</p> <p>To remove an NTP server from the list, select the server to remove, then click the remove button (the X).</p>
apply	Applies date and time configuration changes. (Only visible when accessed through SMU Administration.)
back	Returns to the SMTP Configuration page. (Only visible in SMU Setup Wizard.)
next	Saves configuration changes, and proceeds to the next page of the wizard (private network configuration). (Only visible in SMU Setup Wizard.)
cancel	Returns to the Home page without saving configuration changes. (Only visible in SMU Setup Wizard)

5. When you have specified the date and time settings, click **next** to review the configuration settings you have specified in this wizard.

Field/Item	Description
SMU Settings	Displays user name, but not the password.
Network Settings	Displays the defined DNS servers and search domains.
SMTP	Displays the defined SMTP server.
Date & Time	Displays the currently set date, time, and time zone, and also displays the specified NTP servers.
back	Returns to the Date and Time page.

Field/Item	Description
finish	Saves and implements configuration changes, and proceeds to the Home page.
cancel	Returns to the Home page without saving configuration changes.

- Once you have reviewed all the configuration settings, click **finish** to save the configuration settings and have the SMU start using the new settings.
The SMU will restart, and you must log in using the password specified.



Configuring SMU security


The SMU can be configured to control the hosts that can access the SMU and auxiliary devices managed by the SMU.

Procedure

- Navigate to **Home > SMU Administration > Security Options**.

Field/Item	Description
Control which hosts have access to the SMU	The settings in this section allow you to define the IP addresses of the hosts allowed to access the SMU.
Restrict Access to Allowed Hosts	By filling this check box, you restrict SMU access to only those hosts included in the list of allowed hosts. By leaving this check box empty, you allow any host on your enterprise network to access the SMU.
Allowed Hosts	<p>To allow a host to access to the SMU, enter its IP address here and click add (the down arrow). When first restricting access, this field is pre filled with the IP address of the machine you are currently using to access this page. That IP address is required to be a member of the list. To remove a host from the list of those that have access to the SMU, select the host's IP address in the list, and click delete (the X).</p> <ul style="list-style-type: none"> The format for IPv4 addresses is: <code>#. #. #. #</code>, in which # is a number between 0 and 255. <p>Optionally, you can include a netmask, which is added immediately following the IP address, and is separated from the IP address by a slash (/). The netmask can use either the standard <code>#. #. #. #</code> format, or it can be entered as a simple number between 0 and 32). For example, either of the following are valid: <code>192.168.1.1/255.255.255.0</code> or <code>10.1.1.1/24</code>.</p> <p>The value of specifying a netmask with an IP address is that you can allow access by a range of IP addresses with a single entry. For instance, to allow SMU access only by hosts having an IP address in the range 192.168.1.1 through 192.168.1.255, you could add the single entry <code>192.168.1.1/24</code> instead of entering each of the 255 entries individually.</p>

Field/Item	Description
	 <p>Note: The netmask component does not directly specify the IP address at the end point of a range. For example, entering 192.168.1.1/192.168.1.255 will not allow SMU access for the hosts in the range 192.168.1.1 through 192.168.1.255. Instead, to allow SMU access by all hosts in the range 192.168.1.1 through 192.168.1.255, you would enter 192.168.1.1/255.255.255.0 or 192.168.1.1/24.</p> <hr/> <ul style="list-style-type: none"> The format for IPv6 address is: #: #: #: #: #: #: #: #, for example, fdca:f995:220a:480:1::a (which specifies a single host) or fdca:f995:220a:480:1::a/64 (which specifies a range of IP addresses in CIDR format).
Web Application Security Settings	This section allows you to change web application security settings.
Ports used for SMU access	<p>For added security on your system, you can change the HTTP and HTTPS ports that the SMU uses.</p> <hr/>  <p>Note: If the port settings are changed, the application will be restarted.</p> <hr/>
HTTP	The HTTP port used by the SMU.
HTTPS	The HTTPS (secure HTTP) port used by the SMU.
Enable HTTPS Protocols	By default, all HTTPS protocols are enabled, and the boxes next to the protocols are checked. Uncheck the check box next to a protocol to change its state to disabled. Leave at least one protocol enabled that your browser supports.
Enabled Cipher Suites	By default, all cipher suites are enabled and are shown in the Enabled Cipher Suites list box.
Disabled Cipher Suites	To disable cipher suites, use the arrow to move selected cipher suites to the Disabled Cipher Suites list box. Leave at least one cipher suite enabled that your browser supports.
Login Security Banner	<p>By default, the security banner is disabled. Click Enabled to display the banner on the SMU login screen.</p> <p>The login security banner is displayed on the SMU login screen. The banner file is shared by all login modes (SSH, Serial, GUI, and KVM). A default security banner is provided as a sample security message to users. You can customize this banner text by editing the text on this page.</p> <p>You can also click reset to default, which resets the banner text to the default.</p> <p>You cannot leave the banner empty when creating it using the SMU. However you can leave it empty when creating it using the CLI.</p>
apply	Click apply to save your changes.

Field/Item	Description
	 Note: If change port numbers, clicking apply restarts the SMU, which can take several minutes.

2. Optionally, use the **Restrict Access to Allowed Hosts** check box and the **Allowed Hosts** list to define individual IP address or a range of IP addresses that are allowed to access the SMU and the devices on the private network.
Only hosts from these addresses (or within the defined range of addresses) will be allowed to communicate with the SMU or the devices on the private management network.
3. Optionally, use the **HTTP** and **HTTPS** fields to define the ports that the SMU uses for inbound and outbound communications.
4. Optionally, to disable protocols, at **Enable HTTPS Protocols**, uncheck the check box next to a protocol to change its state to disabled. It is necessary to have at least one protocol remain enabled.



Note: Take care before disabling HTTPS protocols, because not all HTTPS protocols are supported by all browsers.

5. Optionally, to disable cipher suites, use the arrow to move enabled cipher suites from the **Enabled Cipher Suites** list at the left to the **Disabled Cipher Suites** list at the right. It is necessary to have at least one cipher suite remain enabled.



Note: Take care before disabling cipher suites, because not all cipher suites are supported by all browsers.

6. Optionally, click **Enabled**, and change the security banner text.
The security banner is disabled by default.

You can edit the text of the security banner by changing the text in the edit box. Note that the security banner is plain text, and no HTML or formatting is available. To reset the security banner to the default text, click **reset to default**.
7. Click **apply** to save the currently defined security options.

Disabling protocols and cipher suites

Use this procedure to disable individual protocols or cipher suites as by using the **Security Options** page in the SMU Administration menu.

By default, all protocols and cipher suites are enabled. However, occasionally a protocol or cipher suite may be no longer secure and you can use the

Security Options page to prevent a browser from communicating with the SMU using that protocol or suite.

Prerequisites

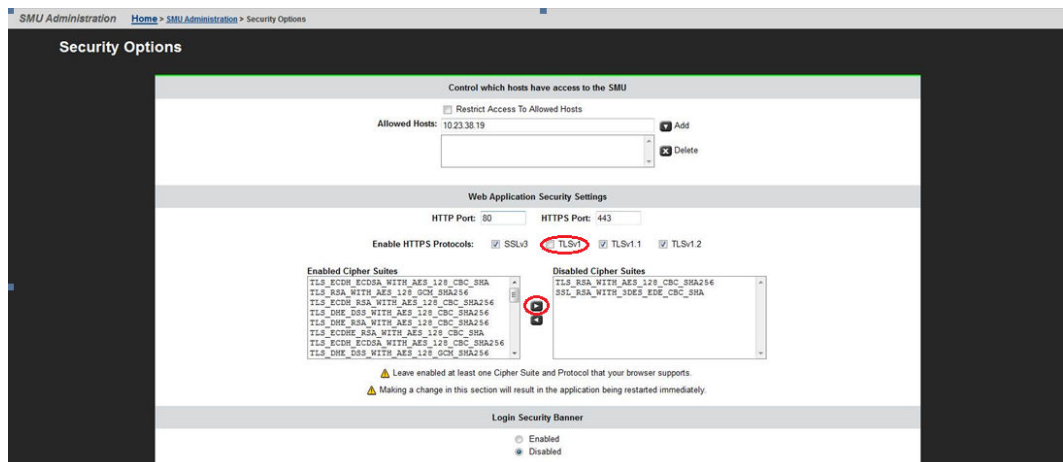
It is necessary to have at least one protocol and cipher suite remain enabled.



Note: It is possible to lose SMU access by disabling too many protocols or cipher suites. If this happens, a warning will appear in your browser. To recover, use the SMU-only CLI command to reset both protocols and cipher suites to defaults. See the man pages for the `smu-reset-tls-options` command.

Procedure

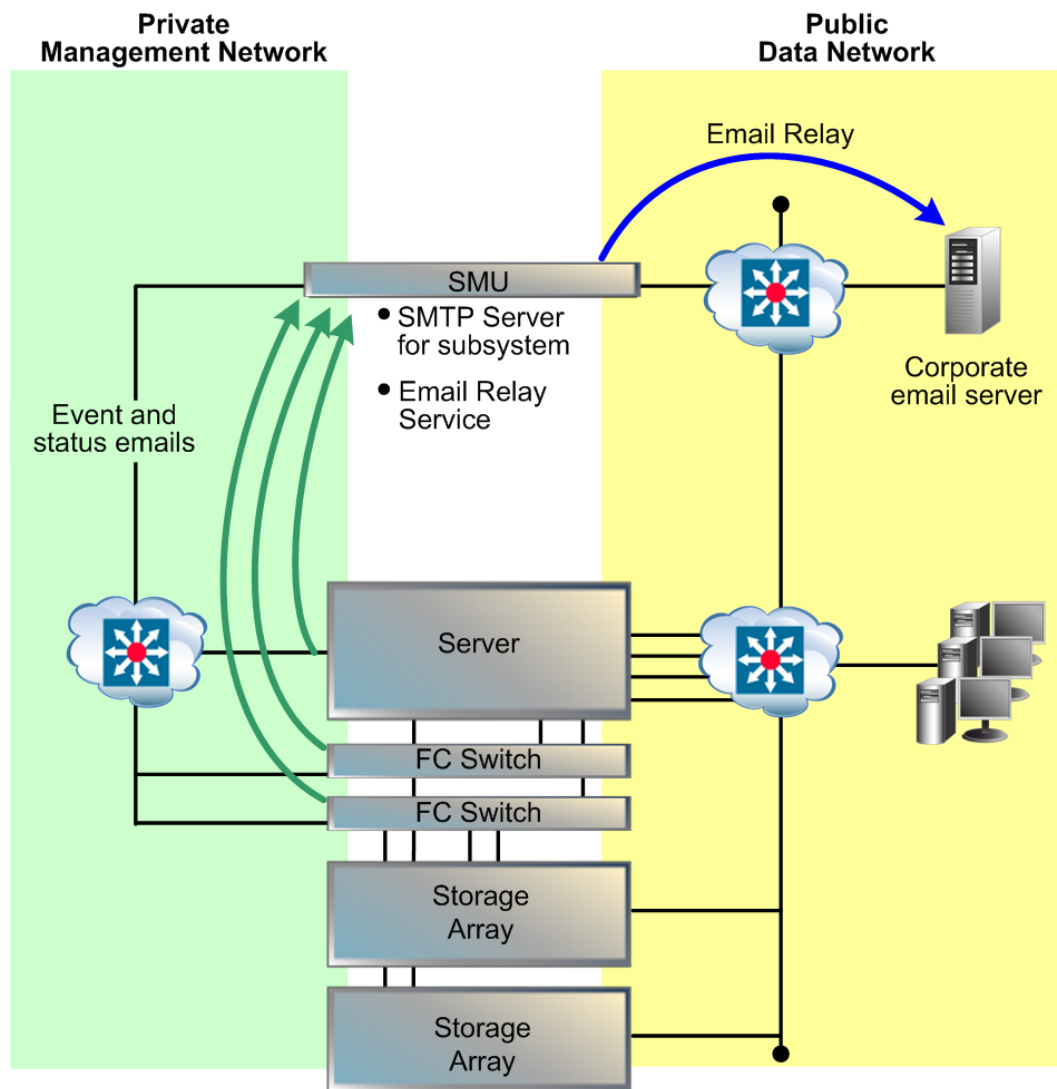
1. Navigate to **Home > SMU Administration > Security Options** to display the **Security Options** page.
By default, all protocols and cipher suites are enabled.
 - a. To disable protocols, at **Enable HTTPS Protocols**, uncheck the check box next to a protocol to change its state to disabled.
 - b. To disable cipher suites, use the arrow to move enabled cipher suites from the **Enabled Cipher Suites** list at the left to the **Disabled Cipher Suites** list at the right.
 - c. Click the **apply** button at the bottom of the page to apply your changes.



Confirmation dialog boxes appear, and then the SMU Application Restarting page. The SMU web application restarts. This may take a couple of minutes.

Configuring an SMTP relay for the SMU

The SMU can be configured to forward emails to the public network from the servers and auxiliary devices on the private management network, using a SMTP relay, as illustrated here:



Procedure

1. Navigate to **Home > SMU Administration > SMTP Configuration**.



Note: If you are accessing this page through the SMU Setup Wizard, the title of the page is **SMU Setup Wizard**, and there are three buttons at the bottom of the page (**back**, **next**, and **cancel**) to take you through the SMU setup.

Field/Item	Description
SMTP Server	Enter the SMTP Server on the public network. The SMU will then relay emails from the servers and other devices on the private network to the public network. An SMTP server can be specified by IPv4 or IPv6 address, or by a

Field/Item	Description
	host name. If an IPv6 address is specified, the SMU will only be able to use the server for email forwarding if the SMU is configured with an IPv6 address. Additionally if the SMTP server is given by host name, and that host name resolves only to an IPv6 address, mail forwarding will only be possible if an IPv6 DNS server is provided.
apply	Saves configuration changes, and closes the page.

2. In the **SMTP Server** field, enter the name of the SMTP server the SMU should use to send emails and through which the SMU will relay email from other devices on the private management network.
3. Click **apply** to save the SMTP server configuration.
4. Verify that the SMTP server IP address specified on the **Email Alert Configuration** page is set to the SMU's eth1 IP address.
View the server's email configuration via the **Email Alerts Setup** link found on the **Status & Monitoring** page.

Displaying the SMU software version

The **About SMU** page displays SMU software version information, including build number and date.

Procedure

1. To display SMU version information, click the **About** link on the bottom of any Web Manager page.
2. Click **back** to return to the page you were on when you clicked the **About** link.

Upgrading SMU software and server firmware

The System Management Unit (SMU) software and storage server firmware can be upgraded to newer releases. For information on upgrading SMU software and firmware, refer to the *System Installation Guide*.

Selecting SMU-managed servers

The SMU manages multiple storage servers/clusters and their associated storage subsystems.

Use the Managed Servers page to add information about each server; specifically, the IP address and username/password of the server to be managed. Only one server, the currently managed server, may be managed at one time. From the Managed Servers list, any server can be selected as the currently managed server.

Procedure

1. Navigate to **Home > SMU Administration > Managed Servers**.

SMU Administration [Home](#) > [SMU Administration](#) > Managed Servers

Managed Servers

▼ IP	Server Username	Model	Cluster Type	Status	
<input type="checkbox"/> 172.31.60.59 - gizmo1	supervisor	Unknown Model	Unknown Type		details Set as Current
<input type="checkbox"/> 192.0.2.3 - g1-cluster	supervisor	3090-G2	Clustered		details Set as Current

[Check All](#) | [Clear All](#)

Actions: [add](#) [remove](#)

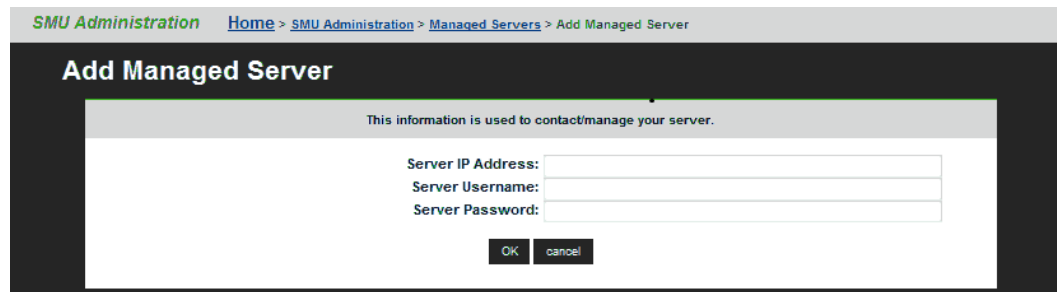
Shortcuts: [Server Upgrade Utility](#) [Server Setup Wizard](#)

Field/Item	Description
IP	IP address of the server. This should be the Administration Services IP address, as used on the private management network (for example, 192.0.2.x).
Server Username	User name of the NAS server.
Model	Displays the NAS server model number. For a cluster with different server models, this field displays "mixed", and the specific server models can be displayed in the Cluster Configuration page.
Cluster Type	Cluster type (for example, Node or Clustered).
Status	The color indicates the current status of the server: <ul style="list-style-type: none"> Green indicates that the server is operating normally (not showing an alert condition). Amber indicates a warning (operating normally, however, action should be taken to maintain normal operation). Red indicates a critical condition (the server is no longer functioning properly).
details	Opens Modify Managed Server page, which contains detailed information about contacting or managing the server.
Set as Current	Makes the currently selected server or cluster the currently managed server/cluster.
add	Adds a server or cluster that will then be managed by this SMU.
remove	Removes one or more selected servers or clusters. When a server or cluster is removed: <ul style="list-style-type: none"> Replication policies and schedules are deleted. Data migration policies and schedules are deleted. The system monitor for that server is deleted. Racks managed by that server are deleted.
Server Upgrade Utility	Opens the Server Upgrade Utility .
Server Setup Wizard	Opens the Server Setup Wizard .

2. Select the servers or clusters the SMU is to manage.

Using the **Managed Servers** page, you can:

- Click **add** to go to the Add Managed Servers page, which you will use to add servers or clusters to the list of managed servers.



Field/Item	Description
Server IP Address	IP address of the server. This should be the Administration Services IP address, as used on the private management network (for example, 192.0.2.x).
Server Username	Username of the NAS server.
Server Password	Password associated with the Server Username.
OK	Saves configuration changes, and closes the page.
cancel	Closes the page without saving configuration changes.

When the SMU adds a managed server, the following actions occur:

- If the server is managed through the private management network, the SMU's eth1 IP address is added to the server's list of NTP servers.
- If the server is managed through the private management network, the SMU's eth1 IP address is configured as the server's Primary SMTP server. If the server was already configured to use a mail server, this server will automatically become the Backup SMTP server.
- A user name and password are preserved on the SMU so that, when using Web Manager, you can select this server as the current managed server without causing the server to prompt for additional authentication.
- Select one or more of the servers or clusters in the Managed Servers list, and click **remove** to delete the server or cluster from the list of managed servers. Make a selection by

filling a server's checkbox, or click **check all** to select all servers in the Managed Servers list.

- Change the *currently managed server*.
Click **Set as Current** to make that server the currently managed server. (Alternatively, you can use the drop-down list in the Server Status console on the **Home** page.)

Changing the IP address of a managed server

If the IP address of a managed server has been changed without using the Web Manager interface (for example, if the server's IP address was changed using the CLI or the console), you can update the IP address used by the SMU to communicate with the managed server.



Note: Updating the IP Address of a managed server does not actually change the IP address of the server, rather it tells the SMU the new IP address of the server. Updating the managed server's IP address does not interrupt management or delete completed replications or data migrations.

Procedure

1. Navigate to **Home > SMU Administration > Managed Servers**.

SMU Administration [Home](#) > [SMU Administration](#) > Managed Servers

Managed Servers

▼ IP	Server Username	Model	Cluster Type	Status	
<input type="checkbox"/> 172.31.60.59 - gizmo1	supervisor	Unknown Model	Unknown Type		details Set as Current
<input type="checkbox"/> 192.0.2.3 - g1-cluster	supervisor	3090-G2	Clustered		details Set as Current

[Check All](#) | [Clear All](#)

Actions: [add](#) [remove](#)

Shortcuts: [Server Upgrade Utility](#) [Server Setup Wizard](#)

Field/Item	Description
IP	IP address of the server. This should be the Administration Services IP address, as used on the private management network (for example, 192.0.2.x).
Server Username	User name of the NAS server.
Model	Displays the NAS server model number. For a cluster with different server models, this field displays "mixed", and the specific server models can be displayed in the Cluster Configuration page.
Cluster Type	Cluster type (for example, Node or Clustered).
Status	The color indicates the current status of the server: <ul style="list-style-type: none"> Green indicates that the server is operating normally (not showing an alert condition). Amber indicates a warning (operating normally, however, action should be taken to maintain normal operation). Red indicates a critical condition (the server is no longer functioning properly).
details	Opens Modify Managed Server page, which contains detailed information about contacting or managing the server.
Set as Current	Makes the currently selected server or cluster the currently managed server/cluster.
add	Adds a server or cluster that will then be managed by this SMU.
remove	Removes one or more selected servers or clusters. When a server or cluster is removed: <ul style="list-style-type: none"> Replication policies and schedules are deleted. Data migration policies and schedules are deleted. The system monitor for that server is deleted. Racks managed by that server are deleted.
Server Upgrade Utility	Opens the Server Upgrade Utility .
Server Setup Wizard	Opens the Server Setup Wizard .

2. Click **details** for the server with the IP address you want to change.

The **Modify Managed Server** page is displayed.

Field/Item	Description
Server IP Address	IP address of the server. This should be the Administration Services IP address, as used on the eth1 port (the 10/100 management port on the private management network). For example, 192.0.2.x.
Server Username	Username of the NAS server.
Server Password	Password associated with the Server Username.
OK	Saves configuration changes, and closes the page.
cancel	Closes the page without saving configuration changes.

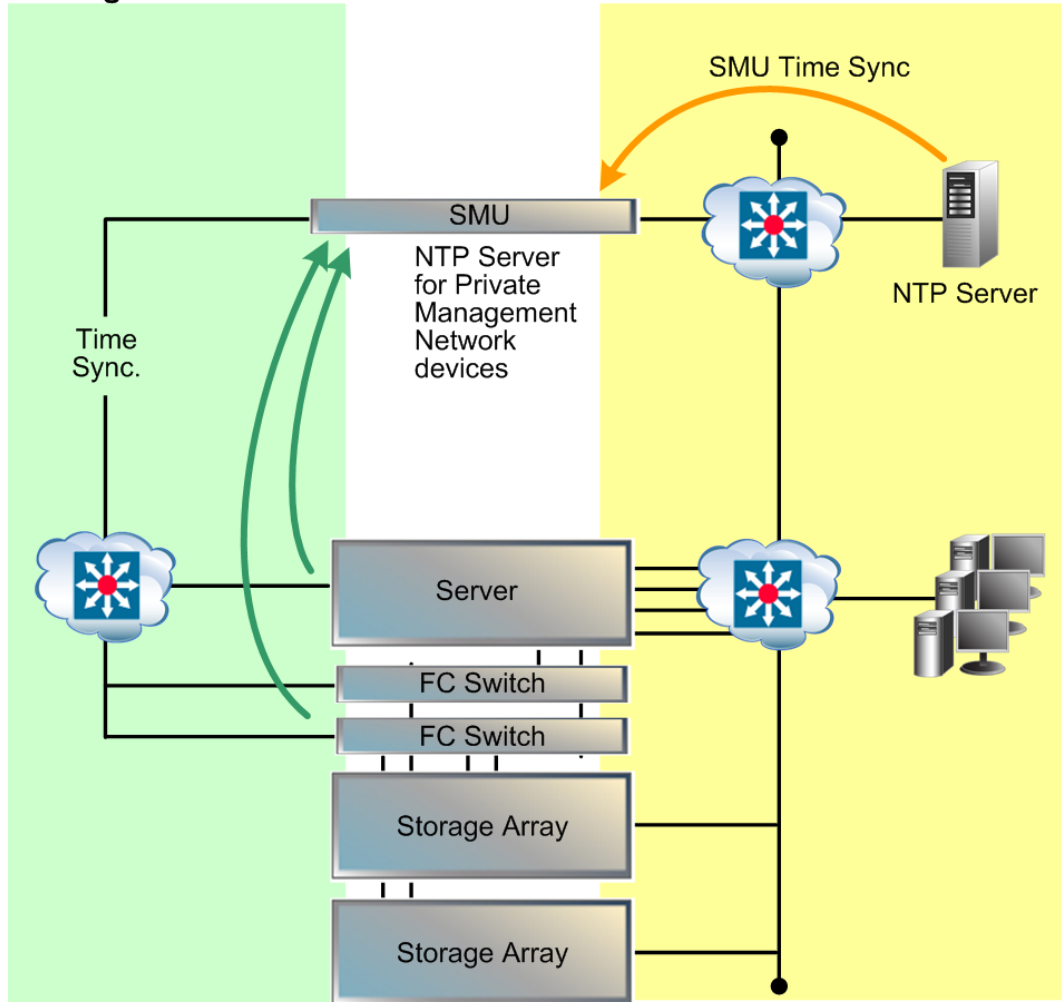
3. Change the IP address as necessary.
4. Click **OK**.
The new IP address is saved, and you are returned to the **Managed Servers** page.

Using the SMU as an NTP server

The SMU is configured as an NTP server. This ensures that every device on the private management network can synchronize with at least one NTP server. In turn, the SMU synchronizes with an NTP server on the public network. The following diagram illustrates this relationship:

**Private
Management Network**

**Public
Data Network**



Configuring the storage server

During the initial setup of the server performed using the **SMU Setup Wizard**, a number of configuration settings, such as system name and date/time were specified. You can import these settings using the procedures in the following the section. Later, you can change these settings, add settings that cannot be imported, and install license keys to enable the protocols and services purchased with the servers.

- ☐ [Cloning server settings](#)
- ☐ [Using the Server Setup Wizard](#)
- ☐ [Configuring server management access](#)
- ☐ [Configuring server identification](#)
- ☐ [Configuring server date and time](#)
- ☐ [Managing license keys](#)
- ☐ [Displaying storage server version information](#)
- ☐ [Upgrading SMU software and server firmware](#)

Cloning server settings

You can copy certain configuration settings to the server. These settings are retrieved from either the SMU or another server managed by the SMU. Depending on the source you select, different settings will be retrieved.

Procedure

1. Navigate to **Home > Server Settings > Clone Server Settings**.

Field/Item	Description
Clone the selected configuration from:	Lists the SMU and the names of all servers or clusters managed by the SMU.
To:	Displays the name of the server to which the settings will be copied.
next	Saves selected source, and proceeds to the next page of the wizard.
cancel	Closes the page without saving the source, and returns to the Home page.

2. Select the source from which you want to copy settings.
 - If this is the first server to be configured, the wizard can clone some settings from the SMU to the new server. Note that the settings that can be cloned from an SMU are a subset of the settings that can be cloned from another server. To clone settings from the SMU, select **SMU** from the drop-down list.
 - If the SMU is already managing another server, an expanded list of settings can be cloned from another server. To clone settings from another server, select one of the managed servers from the drop-down list.
3. Click **next** to display the **Clone Server Settings** page.

Setting can be cloned from:	SMU	Other managed server
Time	X	X
NTP	X	X
Time Zone		X
DNS Servers	X	X
DNS Search Order	X	X
WINS		X
NIS/LDAP		X
NIS/LDAP Servers		X

Setting can be cloned from:	SMU	Other managed server
NS Ordering		X
User Mappings		X
Group Mappings		X
CIFS Domains		X
FTP Configuration		X
SMTP Profiles		X
SMTP Servers	X	X
SNMP Alerts		X
Syslog Alerts		X
SNMP Access		X
Routes		X
NDMP Information		X
Read Cache Options		X
Cloud DM Rules		X
Cloud Accounts and Destination		X

Field/Item	Description
Check All	Click this link to fill the check boxes of all configuration items on the page.
Clear All	Click this link to empty the check boxes of all configuration items on the page.
back	Returns to the previous page of the wizard, where you select the source from which you want to copy settings.
OK	Copies the selected configuration settings to the server, and closes the page.
cancel	Closes the page without saving configuration settings to the server, and returns to the Home page.

4. Select the configuration items you want to clone.
Fill the check box next to each of the configuration items you want to clone. Clear the check box next to the configuration items you do not want to clone
5. Click **next** to clone the settings to the server.
Cloned settings are immediately applied to the server.
6. Reboot or shut down the server.

After you have completed the wizard, you can either reboot the server or shut it down. When the server is restarted, it will use the new configuration.

Using the Server Setup Wizard

This wizard creates a basic server configuration, using user-defined values. At the end of the **Server Setup Wizard**, a confirmation dialog appears, allowing review of settings.



Note: An IP address must be assigned to the server before the Server Setup Wizard can be used. In addition, the server must be added to the SMU as a managed server.

Procedure

1. Navigate to **Home > Server Settings > Server Setup Wizard**.

The pages of the wizard will allow you to enter:

- Server identification, including server description, contact information, and location information.
- IP addresses for the Administrative EVS, Cluster Node 1, and any file serving EVSs on the managed server.
- Name services information, including DNS, WINS, NIS, and name services settings.
- Date and time configuration settings, including time zone and an NTP server.
- CIFS server settings, such as whether the server should use an NT4 or an ADS (Active Directory Service) configuration.
- Email profile information, including SMTP servers, enabling or disabling the built-in support email profile, and email profiles for alerts and system messages.
- Password settings for the Supervisor, Manager, and Root use accounts.

Optionally, you can also create a test file system, CIFS share, and an NFS export.

Once you have gone through all the pages of the wizard, the final **Server Setup Wizard** page displays all the settings that were applied.

2. If necessary, you can run the wizard again to make any changes that are required.

Configuring server management access

The Web Manager provides the primary management interface for managing the server. In certain circumstances, however, an administrator may wish to use one of the following alternatives:

- The command line interface (CLI), accessible through SSH and Telnet.
- The SSC utility, available for both Windows and Linux/UNIX.
- Simple Network Management Protocol (SNMP).

To protect the server from unauthorized access, various safeguards have been built in. Statistics are available to monitor access through these various methods. The following sections detail the configuration options that secure the server's management interfaces and ports.

To prevent unauthorized access to the storage system, you should configure the server to respond only to predefined (authorized) management hosts on the network, based on the management access method (Telnet, SSC and SNMP) and defined port number. You can enable or disable access through SSC and SNMP entirely, and you can specify certain configuration settings to control how those protocols can be used.

Setting the server password

A password is required to authenticate direct management connections to the server. The password is required when adding a server to the SMU's list of managed servers, or when accessing a server directly through the command line interface.

Procedure

1. Navigate to **Home > Server Settings > Change Password**

Field/Item	Description
Current Password	The current password for the currently logged in user.
New Password	The new password for the currently logged in user.
Confirm New Password	The new password again, to confirm the new password for the currently logged in user.
apply	Saves the new password and closes the page.

2. Enter the current password.
3. Enter the new password.
4. Enter the new password again, to confirm.
5. Click **apply** to save the new password.



Configuring SSC access

SSC can be enabled, or disabled, and you can specify the hosts allowed to access the server using this protocol.

Procedure

1. Navigate to **Home > Server Settings > SSC Access Configuration**.

Table 2-1 SSC Access Configuration

Field/Item	Description
Enable SSC Access	Fill the check box to allow access by the SSC protocol, or empty the checkbox to disable access using that protocol.
Port Number	Enter the port number that the storage server should monitor for communication through the protocol. The default is port 206.
Maximum Number Of Connections	Specifies the maximum number of simultaneous connections to the server. You can allow up to five simultaneous connections.
Restrict Access To Allowed Hosts	Fill the check box to restrict protocol access to the hosts specified on this page. Make sure the checkbox is empty to enable the protocol to access any host.
Allowed Hosts	<p>If protocol access is restricted to specified hosts, use these fields to specify the hosts to which the protocol has access.</p> <hr/> <p> Note: If protocol access is restricted to specified to hosts, make sure the SMU is an allowed host.</p> <hr/> <ul style="list-style-type: none"> Allowed Hosts (field). In the Allowed Hosts field, enter the IP address of a host that the protocol is allowed to access, then click Add to insert that host into the list of allowed hosts. <hr/> <p> Note: If the system has been set up to work with a name server, you can identify allowed hosts by IP address or hostname.</p> <hr/>

Field/Item	Description
	<p>Wildcard Usage: You can specify an IP address using the * character, such as: 10.168.*.* or 172.*.*.*.</p> <ul style="list-style-type: none"> Allowed Hosts (list). This list displays the IP address or host name of each of the hosts that the protocol is allowed to access. To delete a host, select its IP address or host name from the list and click Delete.
Add	Inserts that host into the Allowed Hosts list.
Delete	Deletes the selected host from the Allowed Hosts list.
apply	Saves configuration changes, and closes the page.

2. Specify the SSCconfiguration settings.
3. Click apply to save configuration changes.

Configuring SNMP access

SNMP can be enabled, or disabled, and you can specify both the hosts allowed to access the server using this protocol, and which version or versions of SNMP the server will use.

Procedure

1. Navigate to **Home > Server Settings > SNMP Access Configuration**.

Server Settings [Home](#) > [Server Settings](#) > SNMP Access Configuration

SNMP Access Configuration

☐ Send traps upon authentication failure
☐ Disable agent
☐ Process SNMPv1 requests only
☐ Process SNMPv2c requests only
☒ Process SNMPv1 and SNMPv2c requests

Accept SNMP Packets On Port:


Send Traps To Port:


☒ Restrict Access To Allowed Hosts

Allowed Hosts: Add Delete

Allowed Communities: Add Delete

apply

Field/Item	Description
Send traps upon authentication failure	Fill this check box if the SNMP agent is to send a trap in the event of an authentication failure (caused, for example, by the SNMP host using an incorrect community string when formulating a request).
SNMP Protocol Support	Using the radio buttons at the top of the page, select the version of the SNMP protocol with which hosts must comply when sending requests to the agent, or alternatively, disable the SNMP agent.
Accept SNMP Packets On Port	Enter the port number that the server monitors for communication through the SNMP protocol.
Send Traps to Port	Enter the port number that the server uses to send traps.
Restrict Access To Allowed Hosts	Fill this check box to restrict protocol access to the hosts specified on this page. Empty the checkbox to enable the protocol to access any host.
Allowed Hosts	<p>To permit requests from authorized hosts only, type the IP address of a host in this field, then click Add to include it in the list. (If the system has been set up to work with a name server, you can type the name of the SNMP manager host rather than its address.)</p> <hr/> <div>  Note: If access is restricted to specified hosts, add the SMU as an allowed host. </div> <hr/>

Field/Item	Description
	To remove a host from the list, select the host you want to remove, then click Delete .
Allowed Communities	Type the name of a community (a password) that will provide authentication into the MIB, and then click Add to include it in the list. Community names are case-sensitive. <div>  Note: You should define at least one community entry. </div> To remove a community from the list, select the host you want to remove, then click Delete .
apply	Click apply to save configuration changes and close the page.

2. Specify the SNMP configuration settings.
3. Click apply to save configuration changes.

Configuring SNMPv3 access

SNMPv3 defines a more secure version of SNMP compared to the previously supported SNMPv1 and SNMPv2c. SNMPv3 adds support for user-based authentication and encryption to achieve secure access to the management information held on the HNAS server. SNMPv1 and SNMPv2c continue to be available but cannot be enabled at the same time as SNMPv3.

You must use CLI commands to configure SNMPv3.

Prerequisites

The **snmp** concept man page includes information to describe the supported SNMP versions and restrictions.

The authentication and privacy option is always configured when SNMPv3 is enabled.

The SNMP agent uses HMAC-SHA-96 authentication and AES-128-CFB encryption for data privacy.

Procedure

1. Use the CLI command **snmp-protocol** to configure SNMPv3.

```
HNAS1:$          snmp-protocol -v v3

HNAS1:$          snmp-protocol
                  Protocol:      SNMPv3
```

When SNMPv3 is enabled the SNMP agent will not respond to SNMPv1 or SNMPv2c requests.

2. Add users with the **snmpv3-user-add** command.

```
HNAS1:$ snmpv3-user-add testuser
Please enter the authentication password: *****
Please re-enter the authentication password: *****
Please enter the privacy password: *****
Please re-enter the privacy password: *****
[snmpv3-user-add took 14 s.]
```

At least one user, with an authentication password and a privacy password, must be configured in order to use SNMPv3.

When SNMPv3 is configured, access to the information on the server is restricted to users in the SNMPv3 user list.

a. You may delete users with the **snmp3-user-delete** and **snmpv3-user-delete-all** commands

```
HNAS1:$ snmpv3-user-delete testuser
```

b. You may list users with the **snmpv3-user-list** command.

```
HNAS1:$ snmpv3-user-list

      Users
      -----
      testuser
```

3. Configure agent ports using the **snmp-port-set** and **snmp-port-show** commands. The SNMP port used is normally 161.

```
HNAS1:$ snmp-port-set 161
SNMP agent port successfully set to: 161

HNAS1:$ snmp-port-show
SNMP agent port: 161
```

4. The **snmp-trap-port-set**, **snmp-trap-port-show**, and **snmp-traps** commands are available to configure the operation of the SNMP agent for all version of SNMP. The traps are normally sent to port 162.

```
HNAS1:$ snmp-trap-port-set 162

HNAS1:$ snmp-trap-port-show
SNMP trap port: 162
```

All notifications are sent using SNMPv1 traps regardless of the configured SNMP protocol version.

5. When configured to use SNMPv3, the community names configured via the **snmp-communities** command and the hosts list configured via the **snmp-hosts** command do not restrict SNMPv3 access to the server.

Configuring server identification

The server identification information is useful to uniquely identify the server, to provide information about the server's location, and to provide information about who to contact when there are problems.

Procedure

1. Navigate to **Home > Server Settings > Server Identification**.

Server Settings [Home](#) > [Server Settings](#) > Server Identification

Server Identification

Enter descriptive server information.

Server Name:

Description:

Company Name: Identifies this site to your support provider.

Department:

Location:

Address 1:

Address 2: e.g. Hempstead House or Apt 401

City:

ZIP / Postal Code:

State / Province:

Country:

Contact 1:

First Name:

Last Name:

Phone Number:

Email:


Contact 2:

First Name:

Last Name:

Phone Number:

Email:

Field/Item	Description
Server Name	<p>Enter a name for the server (Single Node only).</p> <hr/> <p> Note: When you change the name of the server, the new name will not appear on the Server Status Console (on the Home page) until after the server is restarted.</p> <hr/>
Description	Enter a logical description of the server.
Company Name	Enter the name of the company operating the server.
Department	Enter the name of the department in which the server is installed.
Location	Enter the address of the building in which the server is installed.

Field/Item	Description
Contact 1	Enter the contact information for the person who is the primarily responsible for maintaining the server.
Contact 2	Enter the contact information for the person who is the responsible for maintaining the server when the first contact is not available.
apply	Saves identification settings, and closes the page.

2. Enter the requested information
3. Click **apply** to save the settings.

Configuring server date and time

Administrators configure the server's current date and time, and specify the server's time zone and NTP server for synchronization.



Note: Proper server operation requires time synchronization with a reliable time source. For example, Kerberos authentication (required when operating with Active Directory) depends on the current time. Clock 'drift' may also cause inaccurate reporting of file access and modification times, with unexpected results in data migrations. NTP provides the best and most reliable method for maintaining the server's time accuracy.

Storage server and NTP server interaction

When using NTP, the server first verifies that the specified servers are legitimate; then, over a period of a few hours, gradually adjusts its clock to the time provided by the NTP server. This gradual adjustment is normal, and is designed to minimize the effects of changing the server's clock on utilities that use file timestamps.

If the time initially set on the storage server differs from the time returned by the NTP servers by more than 15 minutes, the server does not try to synchronize to the NTP time; instead, it records a Warning event in the event log, indicating that the date and time must be manually changed to within 15 minutes of the NTP time.

Configuring storage server date and time

Procedure

1. Navigate to **Home > Server Settings > Date and Time**.

Server Settings [Home](#) > [Server Settings](#) > Date and Time

Date and Time

Set the managed server's date and time.

Time: 16:09:08 hh:mm:ss (24 hour)

Date: 2014-06-12

Time Zone: (UTC-08:00) America/Los_Angeles

Set Time at Boot: ☒ If checked, synchronizes time with NTP server(s) during boot.

NTP Server IP/Name:

- 192.168.18.10
- 192.0.2.1
- 192.0.2.2

[Add](#) [Remove](#)

[apply](#)



Note: Proper server operation requires time synchronization with a reliable time source. For example, Kerberos authentication (required when operating with Active Directory) depends on the current time. Clock 'drift' may also cause inaccurate reporting of file access and modification times, with unexpected results in data migrations. NTP provides the best and most reliable method for maintaining the server's time accuracy.

Field/Item	Description
Time	In 24-hour format.
Date	Select from the calendar popup.
Time Zone	Select from the drop-down list. For guidance on which zone to select, see http://www.worldtimeserver.com .
Set Time at Boot	Toggle enabled/disabled to synchronize time with NTP server on reboot: <ul style="list-style-type: none">• If disabled, NTP aligns the server's time with the configured time server gradually and offsets of more than 15 minutes cause NTP updates not to register.• If enabled when the NTP service starts, the time synchronizes immediately, not gradually, and without regard for the current time offset.• A CLI command allows you to connect/disconnect the server from a particular NTP service.
NTP Servers	Enter the IP address of the NTP server(s) you want to use to synchronize the server's time. You can specify several NTP server addresses, and the

Field/Item	Description
	system will qualify and compare all listed NTP servers to determine and set the most accurate time. For servers set up on the private management network, add the SMU's eth1 IP address to the list of NTP servers.
apply	Saves configuration changes, and closes the page.



Note: Never try to compensate for daylight saving by changing the time zone or time in the **Time** and **Date** fields. This can cause synchronization problems if you have a dual server configuration or an NTP server.

2. Set the time, date, and time zone.
3. Specify if the time is to be synchronized with the NTP server at boot time.
4. Specify NTP server or servers for the server to access to get accurate time settings.
5. Click **apply** to save time and date settings, and close the page.

Managing license keys

License keys add powerful services to the storage server, and they can be purchased and added whenever needed. A License Certificate identifies all of the purchased services and should be kept in a safe place. The License Certificate is included in the User Documentation Wallet that was shipped with the system.

System Administrators manage keys for licensed services from the **License Keys** page, which displays the status (and features enabled by) each key and provides controls for adding and deleting keys.

Procedure

1. Navigate to **Home > Server Settings > License Keys** to display the **License Keys** page.

Server Settings [Home](#) > [Server Settings](#) > License Keys

License Keys

MAC ID
MAC ID: d4-28-dd-99-3c-a4

License Key	Cluster	EVS	Storage Capacity	Universal NAS Virtual Model Capacity	Type	Expires	
<input type="checkbox"/> D404-28EC-D341-E8F0-67B9-209D-83B1		0 EVS					details
<input type="checkbox"/> D405-36E1-5A41-E8F0-679B-209D-E487		0 EVS					details
<input type="checkbox"/> D405-B23B-0941-E8F0-67AD-209D-859F		0 EVS					details
<input type="checkbox"/> D406-1628-F741-E8F0-6787-209D-6148		0 EVS					details
<input type="checkbox"/> D406-B041-A841-E8F0-6784-209D-E35F-DE98-1B1C		0 EVS		100 TB			details
<input type="checkbox"/> D406-B059-8E41-E8F0-6784-209D-E35F-DE98-636C		0 EVS		100 TB			details
<input type="checkbox"/> D407-7950-3F41-E8F0-6799-209D-6B6A		0 EVS					details
<input type="checkbox"/> D419-8620-FE41-E8F0-67B6-B68-5B6B-51BD-912D-BFD7	Max 4 Nodes	64 EVS	256 TB				details
<input type="checkbox"/> D41B-521E-1041-E8F0-679C-209D-6AE3		0 EVS					details
<input type="checkbox"/> D41B-7FEE-A441-E8F0-679D-209D-F95F-DE03-B2BE		0 EVS					details

[Check All](#) | [Clear All](#)

Total Licensed on All Unexpired Keys


CIFS	NFS	SFM
WORM	iSCSI	Data Migrator
FS Roll Back	Snapshot Restore	CNS
Read Cache	HDS	DDN
EVS Security	SyncDR	Replication
XVL	FSRS	File Clone
BlueArcRS	Performance Accelerator	Data Migrator Cloud
Premium Deduplication	Extension Pack Secure FTP	

Actions: [add](#) [delete](#) [Show licensed services](#)



Note: The list of licenses presented on your screen may differ from the page shown.

Field/Item	Description
MAC ID	The MAC ID of the server/cluster.
License Key	The alphanumeric string that is the license key.
Cluster	Only displayed for a cluster, this is the maximum number of nodes licensed. This indicates the maximum number of servers that can be configured as nodes of a cluster. Note that cluster licenses are handled somewhat differently than other licenses
EVS	The maximum number EVSs allowed on the server/cluster.

Field/Item	Description
Storage Capacity	The of the maximum amount of storage allowed for the server/cluster, in terabytes. Note that the amount of licensed storage must be equal to or greater than the total amount of storage in all subsystems connected to the storage server or cluster.
Model Type	Where applicable, displays the model type.
Expires	The expiration date for each key, if the key expires (expired license keys are shown in grey).
Total Licensed on All Keys	Displays a list of all services enabled by all installed keys. A checkmark will appear next to the any service enabled by the selected key.
details	To display more details about a particular license key, click details .
Check All	Fills the check box next to each key.
Clear All	Clears the check box next to each key.
add	Displays the License Key Add page.
delete	<p>Fill the check box next to a key, then click delete to remove the key from the server/cluster.</p> <hr/> <div>  <p>Caution: Use extreme care when deleting a license key. Removing the selected license could affect services running on the server.</p> </div> <hr/>
Show licensed services	To see the services licensed by a particular key, fill the check box next to the key and click Show licensed services .

2. Review the current license key information.
3. You can now add or delete license keys.
 - To add a license key, click **add** to go to the **License Key Add** page .
 - To delete a license key, fill the check box next to the key to delete, then click **delete**.

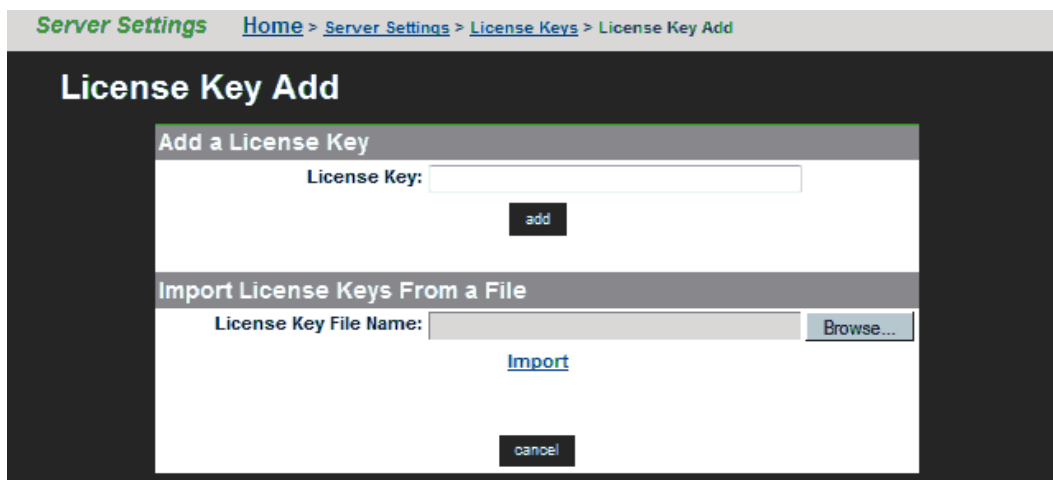
Adding a license key

Adding a license key can enable services or increase the capabilities of your system. To add a license key:

Procedure

1. Navigate to **Home > Server Settings > License Keys**.

2. Click **add**.



The following table describes the fields on this page:

Field/Item	Description
Add a License Key	
License Key	To manually enter the license key, enter the license key and then click add .
Import License Keys From a File	
License Key File Name	To import a license key from a file, click Browse , navigate to the file that contains the license key, and then click Import .
cancel	Closes the page without saving configuration changes.



Note: After adding a license key, if a reboot is required in order to start a service/protocol or enable a feature, you will be instructed to reboot or restart the system.

You can now either enter the key manually or import it from a file:

- To enter the key manually, type it in the field, then click **add**.
- To Import the key, click **Choose File**, navigate to the file, select the key file, then click **Import**.

After all the keys have been entered or imported, they will be displayed on the **License Keys** page. Follow the instructions to reboot the system (if necessary).

License types

Licenses can be grouped into three types:

- Boolean licenses simply enable features/protocols, and when the license is installed the feature/protocol is enabled (for example external volume

links, CIFS, or NFS). These licenses operate in a boolean fashion; if the license is present the feature/function is enabled, if not present, the feature/function is disabled.

- Limit-based licenses specify a limit that cannot be exceeded. These licenses limit your system to a certain total numerical upper limit of the licensed feature/function. For example, the EVS (virtual server) license is a limit-based license.

Limit-based licenses are not cumulative. For example, if your existing cluster has an EVS license for up to nine EVSs, and you install another EVS license for up to eight EVSs, you still cannot have more than nine EVSs (the highest licensed amount). For more information, contact your support representative.



Note: A cluster license is a special kind of limit-based license. When a node joins an existing cluster, its cluster license is transferred to the cluster (if necessary).

- Cumulative licenses. Only the Storage capacity license is cumulative, and several storage capacity licenses can be used to increase the capacity or capability of the system. For example, if you have one storage capacity license for 40 terabytes and another storage capacity license for 60 terabytes, your system could manage up to 100 terabytes of storage.



If a node is removed from a cluster, you must restore its license keys for it to function properly as a standalone server. You should retain the licensing information, in case a node needs to be removed from the cluster.




Note: License keys that have been purchased do not expire. Trial licenses, which enable features for use on a trial basis, have a predefined expiration date. Five days before the expiration of a trial license, the server's event log begins receiving a daily warning event, indicating imminent expiration; then, two days before expiration, the warning events escalate to "severe." When a trial license has expired, the features that enabled by the license become disabled.

The following table lists all services that can be licensed:

Service	Description
CIFS	Common Internet File System. This is a message format used by Windows and MS-DOS to share files, directories, and devices.
Cluster	Clustering. Enables the clustering of up to X nodes. The maximum number of nodes supported in a cluster depends on the series of server used as nodes of the cluster.
CNS	Cluster Name Space. Creates a virtual name space through which multiple file systems can be made accessible using a single mount point. If the EVS Security license is also installed, you can also create individual EVS Name Spaces.

Service	Description
Premium Deduplication	<p>Enables the use of Premium Deduplication, a licensed feature with a four SHA-256 engines, capable of indexing data at a rate of up to 450 MB per second.</p> <hr/> <p> Note: A Base Deduplication feature is automatically enabled by default and does not require a license key. This is a dedupe feature with a single SHA-256 engine (capable of indexing data at a rate of up to 120 MB per second).</p> <hr/>
DM	Data Migrator. Enables more efficient use of primary storage space by transferring older, less performance-critical data to secondary storage.
DMCloud	Data Migrator to Cloud. Enables files hosted on the server to be moved or migrated to public cloud storage.
EVS	Enables up to X Virtual Servers (EVSs). Note that this license is only required for certain NAS server models or clusters using those models as nodes. For those models not requiring this license, the default is to allow the maximum number of EVSs supported by the model.
EVS Security	EVS Security. Enables the creation of Secure Virtual Servers (Secure EVSs).
ExtPackSecureFTP	License for Extension Pack for Secure FTP. A virtual appliance that provides authenticated access to HNAS content via FTP, FTPS, and SFTP.
FileClone	Enables the Writable file clone feature. Refer to the <i>File Services Administration Guide</i> for more information.
FSR	File System Rollback. A tool for restoring a file system to the state of its last successful replication.
FSRS	File System Recovery from Snapshot. A tool for rolling back one or more files in a WFS-2 file system to a previous version without actually copying the data from a snapshot. (For file systems formatted using WFS-2 only.)
HDS	Enables the use of supported storage subsystems manufactured by HDS.
iSCSI	Internet Small Computer System Interface. This license enables iSCSI Initiators to communicate at block level with the servers' iSCSI targets.
Model Type	ModelType. This license upgrades from NAS Platform 4060 to NAS Platform 4080.
NFS	Network File System. This is Sun's distributed file system that enables users of UNIX workstations (including Windows NT systems running an NFS emulation program) to access remote files and directories on a network as if they were local.
QSR	Quick Snapshot Restore. Allows the use of the <code>snapshot-rollback-file</code> and <code>snapshot-recover-file</code> commands, which are used to restore/recover file ss in a file system formatted using the WFS-1 file format. (For file systems formatted using WFS-1 only.)
Read Cache	<p>Cluster Read Caching. Enables Read Caching service, which allows one cached read-only file system per EVS.</p> <hr/> <p> Note: After adding the Read Cache license, you must restart the server/cluster before you can use the read cache.</p> <hr/>

Service	Description
Replication	<p>Replication. Enables replication to external servers (other storage servers, clusters, or an NFS server). This license is required to enable any form of replication outside the server or cluster, including accelerated data copy (ADC). This means that, without a replication license, you can use replication within a server/cluster (you can replicate within an EVS or to a different EVS hosted by the same server/cluster), but you cannot replicate to an external server/cluster.</p> <hr/> <p> Note: The replication license is enforced at the replication source. However, in order to reverse a replication, the source and the target must each have a replication license.</p>
SFM	Server Farm Migration of Virtual Servers. Enables migration of Virtual Servers (EVSs) between servers in a Server Farm.
SyncDR	Synchronous Disaster Recovery Cluster. Indicates if the server/cluster is authorized to use a metropolitan cluster configuration (this is a special configuration, contact Hitachi Data Systems Support Center for more information).
Synchronous Image Backup	SynchImageBackup. This license must be present in order to launch high speed backup
TB	Terabytes. Enables the management of up to "X" terabytes of storage.
UniversalMigratorCapacity (UMC)	(Terabytes) License for <n> TB of virtualization storage. Range is from 1 to 32767 TB of storage, or the value "Max". Max is represented with the numeric value of 32768, but it is to be interpreted as whatever maximum value HNAS supports. This license affects the virtualization of third-party NAS storage
WORM	Write Once Read Many file systems. Used to store crucial company data in an unalterable state for a specific duration.
XVL	External Volume Links. Indicates that cross volume links to data migrated to storage devices attached to a remote server (not necessarily a Hitachi NAS Platform) are enabled. Refer to the <i>Data Migrator Administration Guide</i> for information on cross volume links.

Displaying storage server version information

When requesting technical support, it is important to have version information about storage server firmware and hardware. The following sections explain how to retrieve storage server firmware version information for clusters and stand-alone servers.

Procedure

1. Navigate to **Home > Server Settings > Version Information**.
 - If your system is a cluster, the **Version Information** page is displayed, and it lists the nodes of the cluster along with information about the software version, hardware type, and model number.

Server Settings Home > Server Settings > Version Information				
Version Information				
▼ Cluster Node	Software	Hardware	Model	
Group1-node1	12.1.3600.00	NAS Platform	3090-G2	details
Group1-node2	12.1.3600.00	NAS Platform	3090-G2	details

Field/Item	Description
Cluster Node	Displays the name of the node in the cluster.
Software	Software release for the firmware currently running on the node.
Hardware	Hardware name of the node.
Model	Model number of the node hardware.
details	Displays the Version Details page for the cluster node.

For more information on a node, click the **details** button to view the **Version for Node** page for that node.

- If your system is a standalone-server, the **Version for Node** page is displayed, and this page displays detailed model information of the hardware and version information for the software of the node, including information about the main boards in the server.

Field/Item	Description
Server Version	
Model	The model number of the node hardware.
Software	The software version currently running on the node.
Hardware	The hardware family and the product serial number.
MMBx	Displays the MMB revision. <ul style="list-style-type: none">• mmb: Board release version.
MFBx	Displays the MFB revision. <ul style="list-style-type: none">• mfbxhw : Board type and version information.• Serial Number: MFB serial number.
MCP	Displays the MCP version.

Field/Item	Description
	<ul style="list-style-type: none"> Serial Number: MCP serial number.

Upgrading SMU software and server firmware

The System Management Unit (SMU) software and storage server firmware can be upgraded to newer releases. For information on upgrading SMU software and firmware, refer to the *System Installation Guide*.

Clustering

Administrators can configure a single physical server to act as a standalone server, or as a node in a cluster. The administrator can group several servers into a multi-node cluster or into a server farm.

- A cluster allows multiple physical servers to operate together as a single entity; sharing storage under the centralized management of a single SMU and using a common namespace.
- A server farm allows multiple standalone servers and clusters to be grouped together under the management of a single SMU, sharing a common pool of storage. Each server/cluster in the server farm operates independently of all other servers/clusters in the farm.

File services within the cluster or server farm are virtualized as virtual servers (EVSs), and any file service within the cluster or server farm can reside on, or be migrated to, any node within the cluster/server farm.



Note: When migrating an EVS, both the source and destination server/cluster must be running the same major firmware revision.



Note: A system administrators can query the LDAP server for information about hosts configured into netgroups. You may discover whether a host is in a specific netgroup hierarchy or not, as well as all of the netgroups to which a host belongs. The `nis-is-host-in-netgroup` and `nis-netgroups-for-host` commands are used to check whether a host is a member of a specified netgroup, or to determine the set of netgroups to which a host belongs.

- ☐ [Clusters and server farms](#)
- ☐ [Clusters](#)
- ☐ [Server farms](#)
- ☐ [Clusters versus server farms](#)
- ☐ [Using clusters](#)

- ☐ [Using cluster name space \(CNS\)](#)
- ☐ [Configuring read caching](#)

Clusters and server farms

The key differences between a cluster and a server farm are the behavior when a failure occurs, ease of management, and scalability of operations:

- A cluster allows an EVS to be automatically migrated among cluster nodes in the event of a failure. Management of all nodes in the cluster is centralized, and the usage of a single namespace allows clients to mount a single network resource, while having the actual storage virtualized among different devices attached to the cluster.
- A server farm allows an EVS to be migrated manually among servers in the server farm, but this is a manual process, and it does not happen automatically in the event of a failure. All servers in the server farm may be managed by a single SMU, but each server must be managed as an independent unit.

Clusters

Clustering provides the following functionality:

- Each cluster node can host multiple EVSs. Nodes in a cluster can simultaneously host multiple EVSs, allowing all servers to be active at the same time, each providing file services to clients.
- Redundant monitoring and transparent failover of EVS hosts. The cluster monitors the health of each server through redundant channels. Should one server fail, the EVSs hosted by that node are migrated to another cluster node, which takes over the failed node's functions transparently to network clients, so no loss of service results from a single node failure. After the failed node is restored and is ready for normal operation, previously hosted EVSs can be migrated back.



Note: During the time a node is off line, and during the restoration of the failed node, the cluster may operate with reduced performance.

- Redundant availability of configuration settings for all nodes. The cluster provides a cluster-wide replicated registry, containing configuration information for all nodes in the cluster.

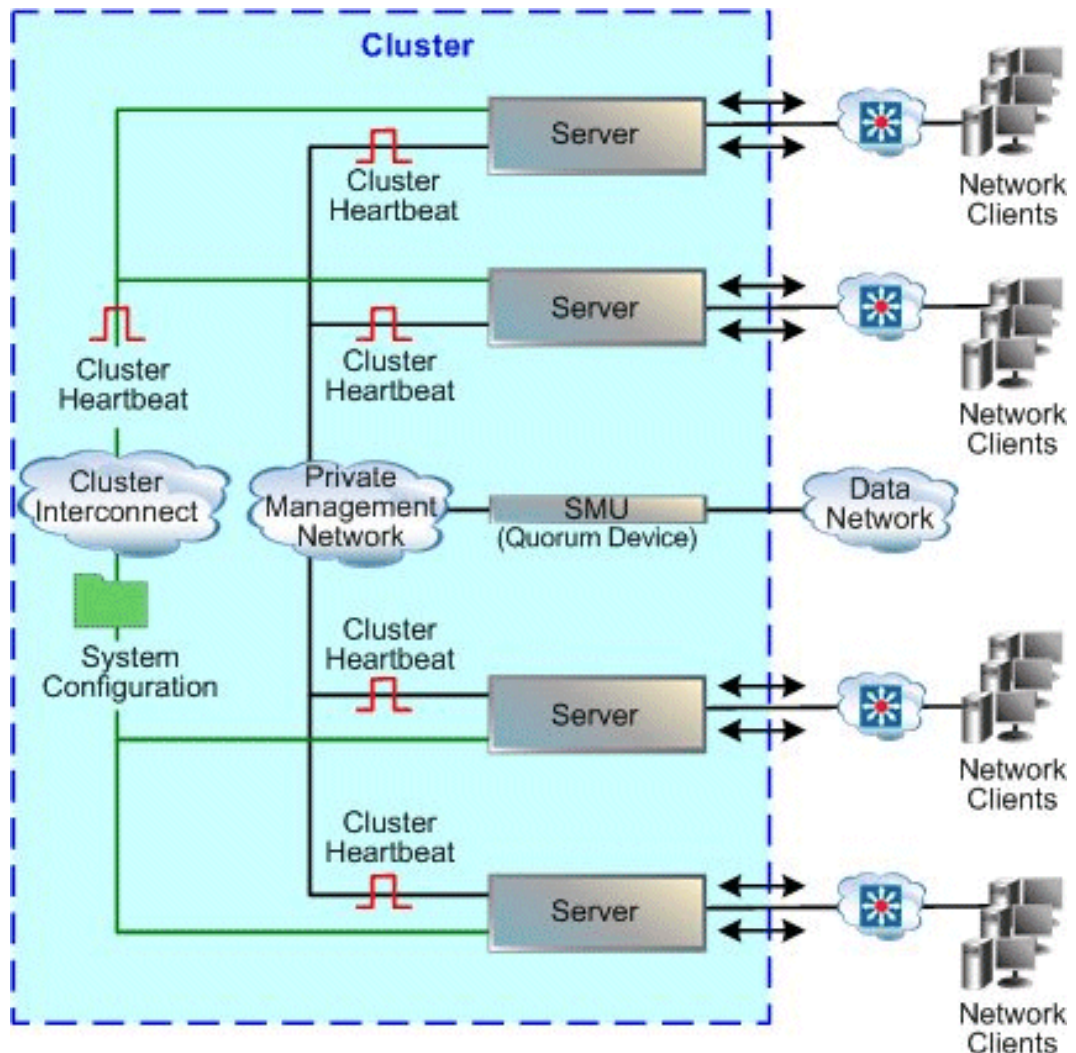
The following sections discuss options for configuring server nodes as clusters, in order to expand their functionality.

Nway clustering

Cluster configurations of more than two (2) nodes are called Nway clusters. The maximum number of cluster nodes is dependent on three factors:

1. The model of server being used as cluster nodes
2. The NAS server firmware version in use
3. The maximum number of cluster nodes allowed by the cluster's licenses.

The following diagram shows the logical view of an Nway cluster configuration of 4 nodes. For more information on setting up Nway clusters, refer to the *Hitachi NAS Platform and Hitachi Unified Storage File Module Series 4000 System Installation Guide*.



Maximum number of nodes supported

The maximum number of nodes in a cluster is controlled by several factors, including hardware version of the server nodes, NAS server software version, and maximum number of cluster nodes allowed by the cluster licenses.



Note: The maximum licensed number of nodes in a cluster will never exceed the maximum number of nodes supported by the hardware and software of the nodes making up the cluster.

For each NAS server model, the maximum supported number of nodes allowed in a single cluster is:

NAS Server Model being used as Nodes	Maximum Number of Nodes Supported
3080	2
3090	4
4040	2
4060	2
4080	4
4100	4



Note: All nodes in a cluster must be of the same model of server.

Quorum device (QD) in a cluster configuration

The quorum device (QD) runs on the system management unit (SMU), which can provide QD services for up to eight clusters (or up to eight servers in a server farm). The QD enables a cluster to maintain operations following a communications failure between nodes and also to restore the cluster registry (containing the cluster configuration), as follows:

- Surviving a communication failure between nodes. Clustering preserves data integrity through a quorum voting algorithm that ensures only one node can access a given file system at any time. Under this algorithm, each of the cluster nodes may “vote” regarding file access. When a cluster contains an even number of nodes, the QD also votes. When a cluster node has obtained a quorum (a simple majority of the votes available in the entire cluster) it receives exclusive access to the file system. Under certain failure scenarios, cluster nodes can lose communication with each other and may attempt to access the same file system; in this situation, the QD alone “votes” for one of the nodes, establishing the quorum and granting one node exclusive access to the file system.
- Preserving a copy of the cluster registry. Although the registry is replicated across cluster nodes, some failure scenarios could result in the loss of recent configuration changes, a condition called amnesia. Anticipating the possibility of such a condition, the QD preserves a copy of the registry, ensuring that configuration changes can always be replicated.

Cluster topology

Typically, the private management network connects cluster nodes and the QD, keeping cluster traffic off of the public data network, and isolating them from potential congestion due to heavy data access loads.

The high-speed cluster interconnect provides an additional, direct connection among the cluster nodes. This dedicated connection consists of dual redundant Gigabit Ethernet (GE) links, and is reserved for clustering traffic and NVRAM mirroring.



Note: Setting up a cluster requires a license. Contact Hitachi Data Systems Support Center to purchase a cluster license.

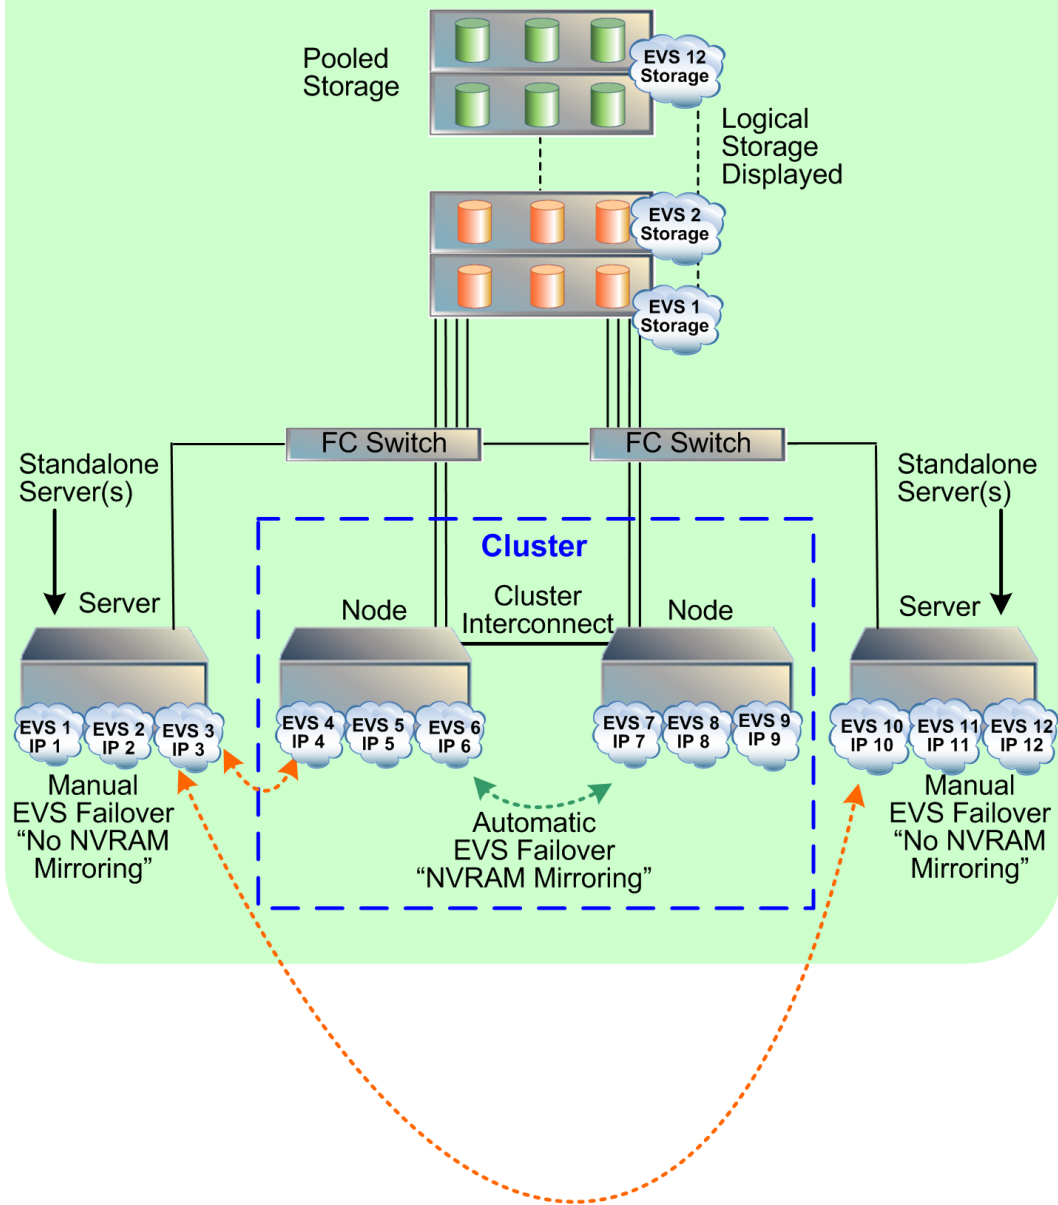
Enhanced cluster quorum device

Release 10.1 includes an Enhanced Cluster Quorum Device, Quorum Services v2. QD v2, hosted on the same system as the SMU, operates in a passive rather than active fashion. It is used only when communication fails between cluster nodes, whereas the previous QD continually polled cluster nodes to detect a node failure. Rather than actively polling the cluster heartbeat, the QD v2 quorum daemon is only called when the server cluster requires an additional quorum vote, in the event of the loss of a cluster node. QD v2 stores node information to elect one node as the master when the cluster experiences a change in the membership. After one node is elected as master, all the remaining nodes will join the master to reconfigure the cluster.

SMU version 10.1 continues to support servers running 8.x firmware, and the SMU runs instances of both the previous and the new Quorum Device daemons. The firmware version determines the user choice of either the legacy QD, or QD v2. Both are managed from the SMU. Servers running firmware 8.x or earlier can only use the legacy quorum services. Servers running firmware 10.x or later, requiring quorum services, can only use Quorum Services v2. The two Quorum Device daemons communicate with each other for cluster configuration, so that a single cluster cannot be served by both the old and the new QD daemons simultaneously.

Server farms

A typical server farm contains at least two standalone servers and/or standalone servers and at least one cluster:

Management
Domain

A single SMU manages every server and cluster within the server farm. The SMU hosts the management network for the server farm and provides quorum services for up to eight clusters. Managed devices must be located in a single data center, not distributed across a campus or MAN environment.

The server farm offers the following functionality:

- **Optimizing performance.** For maximum throughput, migrate EVSs to a higher-end server, or to a fully dedicated server.
- **Balancing load.** For more efficient use of available resources, migrate heavily used EVSs to less-busy servers, or to higher-end servers that support greater capacity.
- **Redundant failover.** In the event of a catastrophic failure of any standalone server, the EVSs hosted by the failed server can be brought online on any other server or cluster in the server farm.

When configured together as a server farm, standalone and cluster nodes share common access to the same storage subsystem, ensuring that when EVSs move from one node to another, whether due to an automatic failover or manual migration of EVSs among servers, the target server has access to the storage served by the EVS.

Clusters versus server farms

The following table distinguishes the properties of a cluster and a server farm:

Property	Cluster	Server farm
Can belong to a server farm	Yes	No
EVS migration under server failure	Automatic	Manual
NVRAM mirroring between servers	Yes	No
Maximum number of storage servers	Depends on several factors. See Maximum number of nodes supported on page 57 for more information.	No explicit restriction on the number of servers; however, an SMU can manage only eight quorum devices and server farm planning should be adjusted accordingly.
Shared SMU	For central management; cluster quorum	For central management; EVS migration
Storage Pools	Yes	No
Common Storage Access	Yes	Yes

Using clusters

Hitachi NAS Platform can form clusters under the following conditions:

- The cluster to which a node is being added must have a license for at least the currently existing number of nodes.
- All nodes in the cluster must have the same hardware configuration. (You cannot form a cluster from a variety of hardware models.)
- The node joining the cluster must be of a compatible software level (within one minor revision level). For example, a server running version 11.0

software can be added to a cluster running version 11.1 software, but not to a cluster running version 11.2 software.

After the first server has been set to cluster mode, you can:

- Add nodes by “joining” servers to the cluster.
- Add EVSs to the cluster and distribute them among the cluster nodes.



Note: In order to maximize cluster performance, distribute EVSs across nodes to level the network client load between them.

Cluster name space (CNS)

A Cluster Name Space (CNS) allows multiple separate file systems on a server to appear as subdirectories of a single logical file system (that is, as one unified file system). They can also make multiple storage elements on that server available to network clients through a single CIFS share or NFS export.

The root directory and subdirectories in the CNS tree are virtual directories. As in a file system, the root occupies the highest position in the CNS tree and subdirectories reside under the root. Access to these virtual directories is read-only. Only the server’s physical file systems support read-write access. Physical file systems can be made accessible under any directory in the CNS tree by creating a file system link. File system links associate the virtual directory in the CNS tree with actual physical file systems.

Any or all of the subdirectories in the CNS can be exported or shared, making them (and the underlying physical file systems), accessible to network clients. Creation and configuration of a CNS can be performed through the Web Manager or the CLI.

After shared or exported, a CNS becomes accessible through any EVS on its server or cluster; therefore, it is not necessary to access a file system through the IP address of its host EVS and, in fact, file systems linked into the CNS can be relocated between EVSs on the server or cluster transparently and without requiring the client to update its network configuration. This can be useful in distributing load across cluster nodes.

The simplest CNS configuration is also the most common. After creating the root directory of the CNS, create a single CIFS share and NFS export on the CNS root; then, add a file system link for each physical file system under the root directory. Through this configuration, all of the server’s storage resources will be accessible to network clients through a single share or export, and each file system will be accessible through its own subdirectory.

Windows and UNIX clients can take full advantage of the storage virtualization provided by CNS, because directories in the virtual name space can be shared and exported directly.



Tip: For the best results, FTP mount points and iSCSI logical units (LUs) should be added to file systems that are not part of a CNS, as CNS does not support FTP mount points or iSCSI LUs. Because FTP clients and iSCSI Initiators communicate directly with individual EVSs and their associated file systems, connectivity for any file system containing FTP mount points or iSCSI LUs must be reestablished through a new EVS upon relocation.



Note: CNS is a licensed feature. To create a Cluster Name Space, a CNS license must be installed. To purchase a CNS license, please contact Hitachi Data Systems.

EVS name spaces

An EVS name space allows separate file systems within a virtual server (EVS) to appear as subdirectories of a single logical file system (that is, as one unified file system). An EVS name space can also make multiple storage elements on the virtual server available to network clients through a single CIFS share or NFS export.

The EVS name space functions in the same way as the cluster name space (CNS), except that its context is that of the EVS, instead of the cluster.

In order to create an EVS name space, you must have installed a CNS license, and an EVS Security license, and you must have set the EVS to use an individual security context.

Linking to and from an EVS name space has the following constraints:

- **Links within an EVS name space.** In an EVS name space tree, you can add links from the EVS name space to file systems hosted by the same secure EVS.
- **Links between the CNS and the EVS name spaces.** The contexts of the Cluster Name Space and the EVS name space are mutually exclusive: links from one to the other are not allowed.
- **Links outside the EVS name space.** Links from the individual EVS name space to file systems in other EVSs are not supported.

About cluster licensing

The maximum number of nodes for a cluster is controlled by several factors, including hardware version, software version, and cluster licenses.



Note: The maximum licensed number of nodes in a cluster will never exceed the maximum number of nodes supported by the hardware and software of the nodes making up the cluster.

A cluster license can be for a single node or for multiple nodes.

- A single node license allows the server/node on which the license is installed to become the first node in a cluster or to join an existing cluster. Using single node cluster licenses, you can form clusters of up to the

maximum number of nodes supported by the hardware/software being used.

Single node cluster licenses can also be used to increase the maximum number of nodes in an already-formed cluster, up to the supported maximum.

- A multi-node license allows the cluster on which the license is installed to form a cluster containing up to the licensed number of nodes, or the supported maximum number of nodes, whichever is lower.

If a server/node containing a multi-node cluster license joins an existing cluster, the cluster's total licensed number of nodes increases to the higher of the following:

- The maximum number of nodes licensed by the existing cluster.
- The maximum number of nodes in the existing cluster's license plus one.

This happens when the total size of the cluster is already greater than or equal to the licensed maximum number of nodes in the existing cluster.



Note: The only difference between a single-node and a multi-node cluster license is the maximum number of nodes the license permits. After installing the license key, you can see the difference between the number of nodes allowed by the license on the License Keys page.

Maximum cluster size can be determined in either of the following ways:

- A cluster containing a multi-node cluster license, for up to "X" nodes.
This method is typically used for new larger-scale installations, where a multi-node cluster is being set up as a new installation and the node containing the multi-node license becomes the first cluster node.
- An additive process, that combines an existing cluster and a node containing a single-node cluster license.

This method is typically used for installations that are expected to grow over time. The key advantage provided by this additive method is that maximum cluster size need not be determined in advance.

For example, you can start with a single server without a cluster license. Later, you install a cluster license, configure the server as the first node of the cluster, and then add nodes. In this situation, you could begin with:

- A multi-node cluster license and then add nodes that don't have cluster licenses into the cluster.
- A single-node cluster license and then install additional nodes (each having their own single-node cluster license) into the cluster.

Another situation where this additive process is used would be if you start with a small cluster, and later add nodes to make a larger cluster. For example, if you start with a two-node cluster that has a four-node license, you can later add two servers (that don't have cluster licenses) to create a four-node cluster. If necessary, you could later grow the cluster by adding

individual nodes (each having a single-node cluster license), up to the supported maximum number of nodes.

Assuming that the cluster has fewer nodes than the maximum size supported by the hardware and software, the rules governing the addition of a node to an existing cluster are fairly simple:

- A node may be added if the licensed maximum number of nodes is greater than or equal to the number of existing nodes, plus one.
- A node may be added if the licensed maximum number of nodes is equal to the number of existing nodes, and the joining node has a cluster license.

When joining an existing cluster, if the joining node has a cluster license, that cluster license is transferred to the existing cluster, and the cluster's maximum number of nodes increases by one (1). The cluster's maximum number of nodes is increased by one, regardless of the maximum number of nodes allowed by the cluster license of the joining node, even if the joining node has a multi-node cluster license. For this reason, the order of joining nodes into a cluster is important.

When becoming a cluster node, all of its licenses are transferred to the cluster, and different licenses are transferred in different fashions.

Configuring a new cluster

Using the Cluster Wizard, you can:

- Create a new cluster by configuring a server as the first cluster node.
- Join a server to an existing cluster as a new node.

Configuring the first cluster node

If any of the nodes that you are going to use to form the cluster contain a multi-node cluster license, that node is the one that should be configured as the first cluster node. Also, the TB (terabytes) license should be installed on the first node of the cluster.

Procedure

1. Navigate to **Home > Server Settings > Cluster Wizard** to display the **Cluster Wizard** page.
2. Enter a new cluster name, associated cluster node IP address, cluster subnet mask, and select a quorum device.



Note: Whether creating a new cluster or joining a cluster node, a cluster node IP address must be defined. This IP address maintains heartbeat communication among cluster nodes and between the cluster nodes and the quorum device (QD), which is typically the SMU. Due to the importance of the heartbeat communication, the cluster node IP address

should be assigned to the 10/100 management port connected to the private management network, keeping the heartbeats isolated from normal network congestion.

3. Click **OK** to save the configuration.
The server reboots automatically. On restart, the node joins the cluster.

Joining an existing cluster using Web Manager

Procedure

1. Navigate to **Home > Server Settings > Join Cluster Wizard** to display the **Join Cluster Wizard** page.
2. Select a server, check the suggested IP address for the node (you can change it, if necessary), enter a user name and password, and click **next**.
3. Allow the system to reboot.
The selected server will automatically reboot, and join the cluster during the boot process.

Configuring the cluster

Procedure

1. Navigate to **Home > Server Settings > Cluster Configuration** to display the **Cluster Configuration** page.

Server Settings [Home](#) > [Server Settings](#) > Cluster Configuration

Cluster Configuration

Cluster Nodes					
Name	IP Address	Status	Model	Health	EVS
Group1-node1	192.0.2.200	Online	3090-G2	OK	
Group1-node2	192.0.2.201	Online	3090-G2	OK	Group1-admin, g1-eva3, donotdelete, g1-eva1, LNAS, g1-eva2, EVS1

Cluster Information

Cluster Name: [rename](#)

Status: Online

Health: Robust

Cluster UUID: a6e6ddfd-9627-11cb-9000-4428dd993ca4

MAC: d4-28-d4-99-3c-a4

Quorum Device

Name: GROUP1-SMU

IP Address: 192.0.2.1

Status: Configured

[add](#) [remove](#)

Actions: [Add Cluster Node](#)

Shortcuts: [Quorum Services v2](#) [EVS Management](#) [EVS Migration](#)

Field/Item	Description
Cluster Nodes	
Name	Node name.
IP Address	IP address of the cluster node. This IP address is on the private management network, which connects devices within the cluster.
Model	Server model, if available.
Health	<p>Worst-case status from each node:</p> <ul style="list-style-type: none"> • OK. Operating normally, with no failures. • Degraded. Operating, but with one or more failures in connectivity. The problem might be with the cluster interconnect, the management network, or quorum device communication. • Failed. Not operating, due to one or more failures in connectivity. The problem might be with the cluster interconnect, the management network, or the quorum device communication. <p>This page also shows the status of operations of the server's internal hard disks, and the percentage of the server's internal disk space that has been used. Disk status is shown as:</p> <ul style="list-style-type: none"> • OK. Operating normally. • Degraded. A non-critical problem has been discovered with one or both of the server's internal hard disks. • Failed. A critical problem has been discovered with one or both of the server's internal hard disks.
EVS	Displays the names (labels) of EVSs hosted on each cluster node. Click the EVS name to display the EVS Details page for that EVS.
Cluster Information	

Field/Item	Description
Cluster Name	Cluster name. Click rename , and enter a new name in the field, to rename the cluster.
Health	Cluster health: <ul style="list-style-type: none"> • Robust. Operating normally, with no failures in the cluster interconnect, the management network, or quorum device communication. • Degraded. Operating, but one or more nodes has failed or there has been a failure in connectivity.
Cluster UUID	UUID (unique ID) of the cluster. This string provides a unique identifier for each cluster when there are several clusters on a network.
MAC	MAC address of the cluster.
Quorum Device	
Name	SMU on which the QD resides.
IP Address	IP address of the SMU on which the QD resides.
Status	QD status: <ul style="list-style-type: none"> • Configured. Attached to the cluster, but vote not needed. The QDs vote is not needed when any cluster contains an odd number of operational nodes. • Owned. Attached to the cluster and owned by a specific cluster node. • Not up. Cannot be contacted. • Seized. Taken over by another cluster.

- As needed, modify the quorum device assignment:
 - Click **add** to assign a QD to the cluster, if a QD is not specified.
 - Click **remove** to remove the specified QD.
If a QD is removed from the cluster, the service will be released back to SMU's pool of available QDs.
- As needed, modify the cluster node assignment:



Note: Services hosted by the cluster node must be migrated to a different cluster node before a node can be removed.

- To remove a cluster node, click its **details** button to display the corresponding **Cluster Node** page. Click **Remove From Cluster**, and **OK** (or **cancel** to decline) in the confirmation dialog.
Upon node removal, any hosted EVSs will automatically be migrated to another cluster node, with details provided in the confirmation dialog.
- To add a node to the cluster, navigate to **Home > Server Settings > Cluster Configuration**, and select **Cluster Join Wizard** to display the **Cluster Wizard**.

Displaying cluster node details

The **Cluster Node Details** page displays information about a selected cluster node and allows removal of that node from the cluster.

Procedure

1. Navigate to **Home > Server Settings > Cluster Configuration**, select a node, and click **details** to display its **Cluster Node Details** page.

Server Settings [Home](#) > [Server Settings](#) > [Cluster Configuration](#) > Cluster Node

Cluster Node Group1-node1

Cluster Node Name: [rename](#)

Cluster Node ID: 1

Status: Online

Network & Storage

File Systems: ● [OK](#)

Ethernet Aggregations: ● [OK](#)

Management Network: ● [OK](#)

Fibre Channel Connections: ● [OK](#)

Cluster Communication

Cluster Interconnect: ● [OK](#)

Management Network: ● [OK](#)

Quorum Device: ● [OK](#)

Chassis

Power Supply Status: ● [OK](#)

Temperature: ● [OK](#) (37 C)

Chassis Disks: ● [OK](#) (Maximum Used: 9 %)

Chassis Battery Status: ● [OK](#) (8.18V)

Fan Speed: ● [OK](#) (3100 rpm)

System Uptime: 2 days 20 hours 20 minutes 23 seconds

EVS

[remove](#)

Field/Item	Description
Cluster Node Name	The cluster node name (label). Click rename to change the name of the cluster node.
Cluster Node ID	The ID assigned to the node.
Network & Storage	
File Systems	<p>Overall indicator of file system status:</p> <ul style="list-style-type: none"> • OK. All file systems up and operational. • Failed. One or more file systems has failed. <p>Click the status link to display the File Systems page, which lists all file systems assigned to the EVS in that cluster node.</p>
Ethernet Aggregations	<p>Overall status of Ethernet aggregations in the cluster node:</p> <ul style="list-style-type: none"> • OK. All aggregated ports are up and linked. • Degraded. One or more ports in an aggregation has failed. • Failed. All ports in an aggregation have failed. <p>Click the status link to display the Link Aggregation page, which lists all aggregations in the cluster node.</p>
Management Network	<p>Overall status of the management network:</p> <ul style="list-style-type: none"> • OK. Links are up and heartbeats are being received. • Failed. No heartbeats are being received, and the links may be up or down. <p>Click the status link to display the Ethernet Statistics page, which lists information about the management port and the aggregated Ethernet ports in the cluster node.</p>

Field/Item	Description
Fibre Channel Connections	<p>An overall status indicator for the Fibre Channel ports in the cluster node:</p> <ul style="list-style-type: none"> • OK. All ports up and operational. • Degraded. Some ports up and operational, but one or more has failed. • Failed. All ports have failed. <p>Click the status link to display the Fibre Channel Statistics Per Port page, which lists all Fibre Channel ports in use in the cluster node.</p>
Cluster Communication	<p>This section contains status indicators for communications within the cluster node.</p> <p>Cluster Interconnect:</p> <ul style="list-style-type: none"> • OK. Link is up and heartbeats are being received. • Standby port down. The primary link is up and heartbeats are being received, but the secondary link is down. • Link up, no heartbeating. At least one link is up, but no heartbeats are being received. • Link down. All links are down (and therefore no heartbeats are being received). <p>Management Network:</p> <ul style="list-style-type: none"> • OK. Both links are up and heartbeats are being received. • Link up, no heartbeating. Both links are up, but no heartbeats are being received. • Link down. Both links are down (and therefore no heartbeats are being received). <p>Quorum Device:</p> <ul style="list-style-type: none"> • OK. The Quorum Device is communicating with the cluster node. • Link up, no quorum communication. The link to the Quorum Device is up, but the Quorum Device is not communicating with the cluster node. • Link down. There is no communication with the Quorum Device.
Chassis	
Power Supply Status	<p>A status indicator for the cluster power supply units (PSUs):</p> <ul style="list-style-type: none"> • OK. Both PSUs are installed and operating normally. • Not Fitted. One PSU not responding to queries, which may mean that has been removed from the chassis, or is not properly installed in the chassis. • Fault or Switched Off. One PSU not responding to queries, and it has failed, been switched off, or is not plugged in to mains power. • Unknown. One PSU not responding to queries, and the exact cause cannot be determined.
Temperature	<p>Status indicator for temperature of the cluster node chassis.</p> <ul style="list-style-type: none"> • OK. Within the normal operating range. • Degraded. Above normal, but not yet critical. • Failed. Critical.

Field/Item	Description
	When available, the temperature in the chassis is also displayed. The displayed temperature is the highest reported temperature of any of the boards in the chassis.
Chassis Disks	<p>Status indicator for operation of the server's internal hard disks, and the percentage of the server's internal disk space that has been used.</p> <ul style="list-style-type: none"> • OK. Operating normally. • Degraded. A non-critical problem has been discovered with one or both of the server's internal hard disks. • Failed. A critical problem has been discovered with one or both of the server's internal hard disks.
Chassis Battery Status	<p>Status of the server's battery pack:</p> <p>When the indicator is green:</p> <ul style="list-style-type: none"> • OK. Capacity and voltage within the normal operating range. • Initialising. PSU battery is initializing after initial installation. • Normal Charging. PSU battery is being charged. • Cell-Testing. PSU battery is being tested. <p>When the indicator is amber:</p> <ul style="list-style-type: none"> • Discharged. Capacity and/or voltage below normal. This status should be considered a warning; if it continues, the PSU battery should be replaced. • Low. Capacity or voltage below normal operating level. This status should be considered a warning; if it continues, the PSU battery should be replaced. • Not Responding. PSU battery is not responding to queries. <p>When the indicator is red:</p> <ul style="list-style-type: none"> • Fault. PSU battery is not holding a charge, has the wrong voltage, or some other fault, and the PSU battery should be replaced. • Not Fitted. PSU battery is not detected. Contact your technical support representative for more information. • Failed. Capacity and voltage consistently below acceptable minimum, or the PSU battery is not charging, or is not responding to queries. This status indicates a failure; the PSU battery should be replaced. • Very Low. Capacity and voltage below acceptable minimum. If this status continues for more than a few hours, it indicates a failure; the PSU battery should be replaced. <p>When available, the level of the battery charge also is displayed.</p>
Fan Speed	<p>Status of fans in the cluster chassis:</p> <ul style="list-style-type: none"> • OK. All fans operating normally. • Degraded. One or more fans spinning below normal range. • Failed. At least one fan has stopped completely, or is not reporting status. <p>When available, the chassis fan speed is also displayed. The displayed fan speed is the slowest reported speed of any of the three fans. An error message may be displayed, even if it does not correspond with the slowest fan.</p>
System Uptime	Duration since last reboot of the cluster node.

Field/Item	Description
EVS	<p>Displays the names (labels) of EVSs assigned to the node, indicates status for each:</p> <ul style="list-style-type: none"> • Green. Online and operational. • Amber. Offline, but listed here because it is hosting the administrative EVS. • Red. Failed. <p>Click the EVS name to display the EVS Details page for that EVS.</p>
remove	Click to remove this node from the cluster.

Quorum device management (external SMUs only)

An external SMU hosts a pool of eight quorum devices (QDs). The SMU provides quorum services for up to eight clusters from its pool of QDs by assigning a QD to a cluster during cluster configuration. After being assigned to a cluster, the QD is “owned” by that cluster and is no longer available for assignment to another cluster. Removing a QD from a cluster releases ownership of the QD and returns the QD to the SMU’s pool of available QDs.

Beginning in SMU software version 10.0, an updated quorum service is available. Depending on version of the NAS server firmware used by the clusters managed by the SMU, one or both of the following quorum service versions may be required:

- Quorum Services (also known as legacy Quorum Services) is required by clusters running firmware versions prior to version 10.0.
- Quorum Services v2 is used by clusters running firmware versions 10.0 and newer.

SMUs running software version 10.0 and later can simultaneously manage clusters that require legacy Quorum Services and other clusters that require Quorum Services v2. Non-managed servers may also use the quorum services on the SMU.



Note: During cluster configuration, the two Quorum Service versions work together to ensure that one cluster cannot be served by QDs of both Quorum Services at the same time. When a request is made to assign a QD to a cluster, the quorum service receiving the request first checks if the other quorum service has already assigned a QD to the cluster. If so, the previously assigned QD is removed from the cluster before the quorum service receiving the request assigns a QD to the cluster. This ensures that the old and new Quorum Services cannot both service the same cluster.

Using cluster name space (CNS)

The CNS has a tree-like directory structure, much like a real file system. Its virtual root and subdirectories provide access to file systems. The CNS can be viewed through the CLI or the Web Manager, and shows all of the configured directories and file system links.

CNS usage considerations

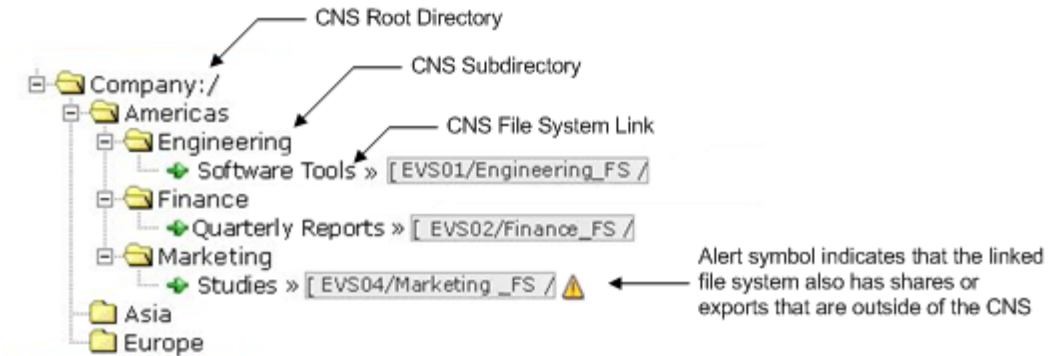
The following recommendations are intended to simplify configuration and maintenance for CNS and for transfers of primary access for the file system:

- A single name space is supported per server or cluster.
- If there is only one CNS link to the file system, and no CIFS shares/NFS exports on the file system, only a single link has to be moved during a transfer of primary access.
- CNS does not support hard links or move operations across the individual file systems. These operations are fully supported, but only within a single physical file system; that is, the part of the CNS tree under a file system link.
- Relocating file systems under the CNS may interrupt CIFS access to the file system being relocated. To minimize interruption, relocate file systems when they are idle. For more information, refer to the *Replication and Disaster Recovery Administration Guide*.
- When using CNS and EVS together:
 - Only one EVS per cluster node is required for all data inside the cluster name space. Having additional EVSs causes unnecessary administrative overhead, and may lead to confusion. Use multiple EVSs on the same cluster node only when you have data that should reside outside the cluster name space.
 - Balance loads by moving file systems, instead of migrating EVS. If you migrate an EVS containing a read cache, the files in the read cache become invalidated and, assuming they are still cacheable, they would have to be cached again after the next read request.
If an EVS containing a read cache is migrated to another cluster node that already has a read cache, the files in the migrated read cache are invalidated, and only the read cache that was not migrated will be used. If the EVS is migrated back to its original cluster node, the read cache will be used again, assuming another read cache has not been created on that cluster node in the interim.
- When using CNS, the recommended configuration is to have a single CIFS share or NFS export at the root of the name space. If that configuration does not suit your needs, the next-best configuration is to have CIFS shares/NFS exports pointing to individual directories in the name space. You should not configure CIFS shares or NFS exports pointing to a path of the real file system unless absolutely necessary.

Displaying the cluster name space tree

Procedure

1. Navigate to **Home > Files Services > CNS** to display the **CNS** page.

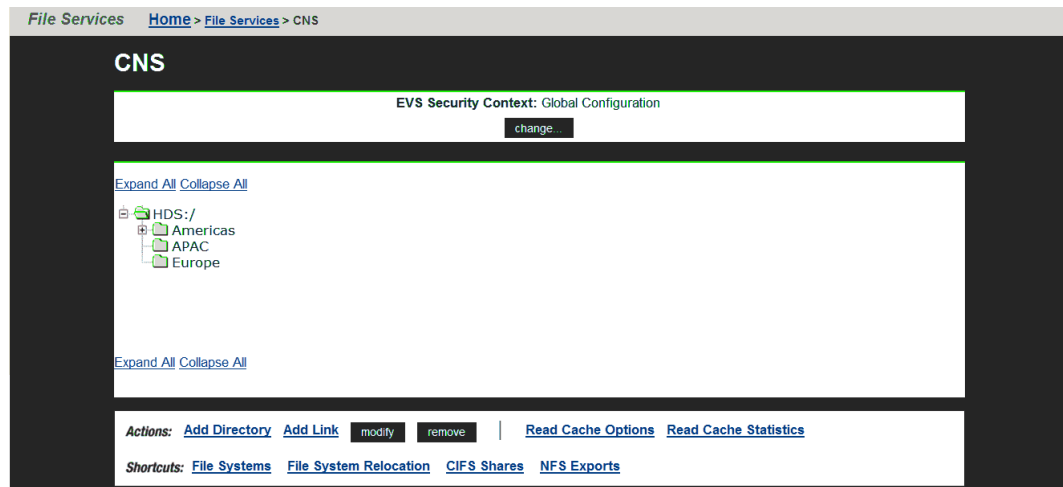


2. If a secure EVS has been created and you want to display the EVS name space for that secure EVS, click **change** to select the name space you want to display.
After you have selected a name space, the tree for that context is displayed.
 - At the top of the name space is the root directory.
 - Under the root directory are a number of subdirectories. In this example topology, one subdirectory has been created for each physical file system.
 - Under each subdirectory is a file system link. A file system link associates a directory with a specific file system. The EVS to which the file system is associated is also shown.

Displaying the EVS name space tree

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page. The currently selected EVS security context and the current name space are displayed at the top of the page, and the tree for the current name space is displayed below the name space label.



2. Click **change** to display a list of name spaces (the Global Configuration, and all individual EVS name spaces that have been defined).
3. Click **Global Configuration** or the EVS name space to display the tree for that name space.
 - At the top of the name space is the root directory.
 - Under the root directory are a number of subdirectories. In this example topology, one subdirectory has been created for each physical file system.
 - Under each subdirectory is a file system link. A file system link associates a directory with a specific file system. The EVS to which the file system is associated is also displayed.

Managing links and subdirectories in the EVS name space

Links and subdirectories in an individual EVS name space are managed in the same way as they are in the CNS.

Creating a cluster name space tree

A CNS contains a root directory, file system links, and, optionally, subdirectories. The first step required to configure CNS is to create the root directory.

Creating a CNS root directory

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. In the **CNS Root Label** text box, enter a name, and click **OK** to create the CNS.



Note: For the CNS to be available to clients, a CIFS share or an NFS export must be created for it. See the *File Services Administration Guide*.

Creating CNS subdirectories

Subdirectories can be created under the root directory or under other subdirectories in the CNS tree. They are optional, but give structure to the CNS, allowing granular control over the organization of physical file system resources.

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. Click **Add Directory** to display the **Add CNS Directory** page.

3. From the **Select a Parent for the Directory** options box, select a parent directory, and enter a name in the **Subdirectory Name** text box.
4. Click **OK** to create the subdirectory, and repeat to add any additional subdirectories.

Creating a file system link

File system links make physical file systems accessible through the CNS. A file system link can be associated with either the root directory or a subdirectory in a physical file system. After created, a file system link is displayed as a directory in the CNS. The directory name seen by a network client will be the name given to the file system link. A network client

navigating through CNS and into a file system link will see the contents of the directory that was linked.

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. Click **Add Link** to display the **Link File System** page.

3. In the **Link Name** text box, enter a name for the link.
4. In the **From CNS Directory** options box, select a location in the CNS tree to place the link.
5. To allow files or cross file system links from remote file systems to be read cached, go to the **Remote Read Cache** list, and select one of the following:



Note: For information about read caching, see the *File Services Administration Guide*.

- **Cache all files.** Allows caching of files from a file system hosted by an EVS on a remote cluster node, and files accessed by local links to a remote file system (cross file system links). A remote cluster node is a node other than the one to which the client is connected.
- **Cache cross file system links.** Allows only cross file system links to be cached.



Note: When the link being added is for a file system in an EVS that has an EVS individual namespace, remote read caching is not available.

To disallow read caching of files from remote file systems, do not change the default selection of **Do not cache files**.

6. To allow files or cross file system links from local file systems to be read cached, go to the **Local Read Cache** list, and select one of the following:
 - **Cache all files.** Allows caching of files from file systems on the same server/node as the read cache, and files accessed by local links to a remote file system (cross file system links). The remote file system might be a remote server or storage device.
 - **Cache cross file system links.** Allows only cross file system links to be cached.

To disallow read caching of files from remote file systems, do not change the default selection of **Do not cache files**.

Changing cluster name space properties

After a CNS has been created, any of its properties can be changed, except the name of the root directory.

Deleting a cluster name space

Deleting a CNS will permanently erase it. Deleting a CNS will not affect the physical file systems accessible through the CNS. However, once the CNS has been deleted, it may be necessary to restore access to the file system by sharing or exporting the file system through its EVS.

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. From the CNS directory tree, select the CNS root directory, and click **remove** to open a confirmation dialog.
3. Click **OK** to delete the CNS.

Renaming a CNS subdirectory

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. From the CNS directory tree, select the subdirectory to be renamed, and click **modify** to display the **Modify CNS Directory** page.
3. In the **Subdirectory Name** text box, enter a new name for the CNS directory, and click **apply** to open a confirmation message box.
4. Click **OK** to rename the directory.

Moving a CNS directory

Moving a CNS directory from one location in the CNS to another can be done at any time.

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.

2. From the CNS directory tree, select the subdirectory to be moved, and click **modify** to display the **Modify CNS Directory** page.
3. From the **Select a Parent for the Directory** options box, select a new location in the CNS tree. From the bottom of the options box, click **apply** to open a confirmation message box.
4. Click **OK** to move the directory.

Deleting a CNS directory

Deleting a CNS directory will permanently remove it and all of its subdirectories and file system links. Deleting CNS directories will not affect physical file systems on the server.

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. From the CNS directory tree, select a subdirectory.
3. From the box at the bottom of the page, click **remove** to open a confirmation message box.
4. Verify your settings, and click **OK** to proceed, or **cancel** to decline.

Modifying a file system link

The name and location of a CNS file system link can be modified.

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. From the CNS tree, select a file system link, and click **modify** to display the **Modify Link** page.
3. As needed, change the link name or parent directory.
 - To change the name of the file system link, enter the new name in the **Link Name** text box, and click **apply**.
 - To change the parent directory, select a new location in the tree from the **Select a New Parent Directory** options, and click **apply**.
4. If necessary, change the setting to enable or disable the caching of files from this file system, and click **apply**.

Deleting a file system link

Procedure

1. Navigate to **Home > File Services > CNS** to display the **CNS** page.
2. From the CNS tree, select a link, and click **remove** to display a confirmation dialog.
3. Click **OK** to proceed.

Configuring read caching

Prerequisites

A storage server can support read caching under the following conditions:

- License keys to enable the read caching service and the Network File System (NFS) service must be installed.
- Sufficient space must be available in a storage pool to create the read cache.

Additionally, to support remote read caching:

- The storage server must be configured as a part of cluster.
- The cluster name space (CNS) feature must be properly licensed and configured.



Note: After the read cache license key is entered, the server/cluster must be restarted before the read caching service starts.

Before you can configure the read caching service, you must have already fulfilled the prerequisites.

To enable and configure read caching, you must:

Procedure

1. Enable the read caching service.

To enable read caching, you must add the license key for read caching. After the key has been added, the service will be enabled upon restart, enabling creation of read caches.

2. Create a read cache on an EVS.

Because a read cache is a kind of file system, the same procedure that creates file systems also creates read caches. For information about creating a read cache or a file system, refer to the *File Services Administration Guide*.

3. Enable file caching.

The configuration can specify that files from some file systems should be cached, while prohibiting file caching for files from other file systems. To control the caching of files from a file system, select the file caching option when you add the file system link or export.

4. Set file caching options.

To control which files are eligible for caching, you must configure the file caching options. After a file system link has been added to the CNS tree, the file system link options can be changed to control whether files from this file system can be cached.

Configuring file caching options

Procedure

1. Navigate to **Home > Storage Management > Read Cache Options** to display the **Read Cache Options** page.

[Storage Management](#) [Home](#) > [Storage Management](#) > [Read Cache Options](#)

Read Cache Options

Cache files on specified NFS exports or CNS links which conform to the following characteristics

Before caching, a file must remain unmodified for: Minutes ▼

Do not cache files larger than: MB ▼


Files may be removed from cache if not accessed for: Minutes ▼

Once removed, a file cannot be cached again for at least: Minutes ▼

Actions: [Reset Values to Default](#)

Shortcuts: [Read Cache Statistics](#)

Field/Item	Description
Before caching, a file must remain unmodified for	<p>Specifies how long a file must be unchanged before it is eligible for caching.</p> <p>The default minimum stable time is 10 minutes.</p> <p>This does not indicate how long since the file has been accessed, only that the file may not have been changed within this period. You can specify the time in seconds, minutes, hours, or days.</p>
Do not cache files larger than	<p>Limits the maximum size of a file eligible for caching.</p> <p>The default maximum size is 512 megabytes.</p> <p>Caching large files might limit the number of files that the read cache can contain. If necessary, you can expand the read cache as described in the <i>File Services Administration Guide</i>.</p>
Files may be removed from cache if not accessed for	<p>Which indicates the amount of time that a file will remain in the read cache without being accessed before it is designated as inactive. Inactive files are eligible for removal from the cache, and are deleted on an "as space is needed" basis, with the oldest inactive files being deleted until there is enough space for a new file to be added to the read cache.</p>

Field/Item	Description
	The default duration is 15 minutes. You can specify the amount of time in seconds, minutes, hours, or days.
Once removed, a file cannot be cached again for at least	<p>Specifies the minimum time that must elapse before a file is re-evaluated for read caching after:</p> <ul style="list-style-type: none"> • Having been read cached, and then having been flushed from the read cache for any reason. • Being evaluated for read caching, and not being cached for some reason. Reasons for not caching a file may include file size, too recent modification, or insufficient space available in the read cache. <p>The default retry time is 30 minutes. You can specify the amount of time that must elapse before a file can be cached again in seconds, minutes, hours, or days.</p>
apply	Click to save the values specified for the options on this page.
Reset Values to Default	<p>Click to reset the values on this page to their defaults.</p> <hr/> <p> Note: This link does not reset the active file set, which must be done through the CLI (Command Line Interface).</p> <hr/>
Read Cache Statistics	Click this shortcut to go to the Read Cache Statistics page.

2. Set the options, and click **apply** to apply changes, or **Reset Values to Default** to restore defaults.



Note: You can prohibit read caching of files from a particular file system when the link to that file system is added to the CNS tree.

Reviewing read cache statistics

Read cache statistics provide information about a read cache, including:

- **Successfully Cached Files:** The number of successfully read cached files.
- **Candidate Files Encountered:** For remote read caching, the number of read cacheable files that have been read by a remote node. For local read caching, the number of read cacheable files that have been read by the local node.
- **Files Rejected:** Has Named Streams: The number of read cacheable files that were not cached because they have associated named streams.
- **Files Rejected: Not Stable:** The number of read-cacheable files were not cached because they were modified within the window of time specified by the "Before caching, a file must remain unmodified for" setting.

- **Files Rejected: Too Large:** The number of read-cacheable files that were not cached because they exceed the size specified in the “Do not cache files larger than” setting.
- **Flushes Due To Active Set Limit:** The number of times a file was flushed from the read cache because the read cache reached its maximum number of active files. By default, a maximum of 250,000 files may be in the read cache at any one time.
- **Flushes To Reclaim Space:** The number of times a file was flushed from the read cache to free space in the read cache file system.
- **Flushes Aborted:** The number of times an unaccessed file in the read cache was not flushed because it was still considered active according to the “Files may be removed from cache if not accessed for” setting.
- **File Lock Revoked:** The number of times all files in the read cache were invalidated at the same time. This statistic also counts the times a file lock is revoked because the “real” file has been modified, which causes the cached copy to be removed from the read cache.

Certain situations will cause the simultaneous invalidation of files in the read cache; some are the result of normal operations (like the unmounting of a file system), while others are due to error conditions.

For local read caching, this situation occurs whenever the local file system is unmounted, or when the EVS hosting the file system is migrated to another cluster node.

For remote read caching, this situation occurs whenever there is a loss of communication with the remote file system. The number of times all files in the read cache from a particular remote file system were invalidated at the same time. This situation occurs whenever there is a loss of communication with the remote file system.

For example, all files from a particular remote file system are invalidated simultaneously when:

- The remote file system is unmounted.
- The cluster node on which the remote file system is located crashes.
- The cluster interconnect fails.
- **Average Cached File Size:** The average size of files stored in the read cache.
- **Average Cached File Lifetime:** The average time a read cached file stays valid and can therefore service read requests. This statistic provides a very good indicator of the efficiency of the read cache. If this value is small, it can have several causes:
 - Files are being flushed from the cache too often because too many files are being cached.

If files are being flushed from the cache too often, consider reducing the number of CNS links marked as read cacheable, or increasing the maximum number of files allowed in the read cache. (Contact Hitachi Data Systems Support Center for assistance if you want to increase the maximum number of files allowed in the read cache.)

- Files are being flushed from the cache because the read cache is running out of space.
If the read cache is running out of space, you can increase the size of the read cache file system, or you can decrease the number of files that are cached (either by decreasing the maximum number of files allowed in the read cache, or by reducing the number of CNS links marked as read cacheable).
- Files that are identified as read cacheable are actually being modified too often.
If files identified as read cacheable are being modified too often, increase the value of the “Before caching, a file must remain unmodified for” option.

Displaying read cache statistics

Procedure

1. Navigate to **Home > Status & Monitoring > Read Cache Statistics** to display the **Read Cache Statistics** page.

Status & Monitoring [Home](#) > [Status & Monitoring](#) > Read Cache Statistics

Read Cache Statistics

▼ <u>Cluster Node</u>	<u>Active Read Cache</u>	
<input type="checkbox"/> Group1-node1	-	details
<input type="checkbox"/> Group1-node2	-	details

[Check All](#) | [Clear All](#)

Actions: [Reset](#)

Shortcuts: [Read Cache Options](#)

2. Select a read cache, and click **details** to display its statistics page.

Status & Monitoring [Home](#) > [Status & Monitoring](#) > [Read Cache Statistics](#) > Read Cache Statistics

Read Cache Statistics on Cluster Node Group1-node1

File System	Number of Cached Files	Total Size Cached
There are no file systems caching files on this cluster node.		

Successfully Cached Files: 0
Candidate Files Encountered: 0
Files Rejected: Has Named Streams: 0
Files Rejected: Not Stable: 0
Files Rejected: Too Large: 0
Flushes Due To Active Set Limit: 0
Flushes To Reclaim Space: 0
Flushes Aborted: 0
File Lock Revoked: 0
Average Cached File Size: 0 Bytes
Average Cached File Lifetime: 0.00 Milliseconds

[back](#) [reset](#)



Note: The table at the top of the **Read Cache Statistics** page lists the name of each file system that currently has files in the read cache. For each file system that currently has cached files, the table lists the number of files and their total (aggregated) size.

3. After you have reviewed the available statistics, you can:
 - Click **reset** to restart the gathering of statistics (you will lose previously gathered statistics for the read cache).
 - Click the **back** button to return to the **Read Cache Statistics** page.

Deleting a read cache

For information on how to completely delete a read cache, refer to the *File Services Administration Guide*.

Read cache considerations

The following recommendations are intended to take full advantage of read caching:

- Because remote read caching requires CNS, you should review the cluster name space considerations.
- In a cluster configuration, define one EVS per cluster node, and assign a read cache to each EVS.
- Balance loads by moving file systems, instead of migrating EVS. If you migrate an EVS containing a read cache, the files in the read cache become invalidated and, assuming they are still cacheable, they would have to be cached again after the next read request.

If an EVS containing a read cache is migrated to another cluster node that already has a read cache, the files in the migrated read cache are

invalidated, and only the read cache that was not migrated will be used. If the EVS is migrated back to its original cluster node, the read cache will be used again, assuming another read cache has not been created on that cluster node in the interim.

- Do not relocate read caches. If you relocate a read cache, the files in the read cache become invalidated and, assuming they are still cacheable, the previously cached files would have to be cached again after the next read request.

Using virtual servers (EVSs)

A server node supports up to 64 EVSs. Likewise, a cluster can have up to 64 EVSs. EVSs can be added, deleted, and changed based on the evolving needs of the network.

- ☐ [Secure virtual servers](#)
- ☐ [Secure EVS considerations](#)
- ☐ [Securing an EVS](#)
- ☐ [Removing an individual security context from a secure EVS](#)
- ☐ [EVS name spaces](#)
- ☐ [Creating an EVS](#)
- ☐ [Assigning a file system to an EVS](#)
- ☐ [Virtual server \(EVS\) management](#)
- ☐ [Displaying EVS details](#)
- ☐ [Migrating an EVS within a cluster](#)
- ☐ [Migrating an EVS within a server farm](#)

Secure virtual servers

A secure virtual server is a file serving EVS that has a specifically defined security configuration (called an individual security context). When no individual security context is specified for an EVS, it uses the global (server or cluster-wide) security configuration settings (the global security context). By defining an individual security context for a particular EVS, you create a secure virtual server (secure EVS).



Note: Secure virtual servers are a licensed feature, identified as EVS Security. Without an EVS Security license, all EVSs use the global security settings (context).

- When no individual security context is defined for an EVS, the global security settings (the global context) are used by default. When an individual security context is added to an EVS, the new individual security context is created using the same settings as are used by the global security context. After adding the individual security context, you can then change settings to make the individual security context settings different than the global settings.
- When using an individual security context, the EVS security context can be configured independently of the global (server or cluster-wide) security settings.

When present, individual security context settings override the global security context settings, allowing a storage server (or cluster) to be shared by multiple groups (departments, customers, or organizations), while maintaining strong security so that no group has access to another group's data.

For example, if a server/cluster has six EVSs, you could define individual security contexts for two of the EVSs, turning them into secure EVSs. Each secure EVS could then be associated with an NT domain that is different than the one used by the cluster, meaning that each of those secure EVSs could be assigned to its own domain. For network clients, access to the file systems in the secure EVSs can then be restricted or allowed as desired using standard network security policies such as user name or user group membership.

Secure EVS considerations

When using secure EVSs, keep the following points in mind:

- **Security context defaults.** Unless an individual security context is specified for an EVS (making it a secure EVS), the EVS security context defaults to the global security context.

- **Inherited global settings.** NDMP user name and password settings are not EVS-specific; the same NDMP user name and password settings apply to all EVSs and secure EVSs in a server/cluster.
- **Secure EVS-specific security settings.** After an EVS has a defined individual security context, it becomes a secure EVS, and each secure EVS is considered to be separate from all other EVSs and secure EVSs in the server or cluster.

A secure EVS is always treated as an individual unit, regardless of if it uses the same security context settings as another secure EVS or if it uses different security context settings. As a result, different secure EVSs cannot share anything, including an individual EVS name space.

- **Secure EVS migration.** When a secure EVS migrates to a different cluster, it retains all specified security settings in its individual security context. If, however, a secure EVS is configured to use default settings from the global context, after migrating, the secure EVS switches to use the settings in the global context of the cluster to which it migrates.
- **Moving file systems between secure EVSs.** A system administrator with sufficient privileges can move a file system from one secure EVS to another, but a warning is issued if the security contexts of the source and destination secure EVSs are different.
- **External name server access.** Each secure EVS can be configured to connect to several external name servers, and each secure EVS can connect to different name servers.
- **Secure EVSs and name spaces.** Links from the cluster's CNS tree to a secure EVS are supported, according to the following rules:
 - **Accessing the CNS.** Only a secure EVS that uses the global security context can access links in the CNS.
 - **CNS links to a file system hosted by a secure EVS with an individual security context are not allowed.** In the CNS, you cannot add a link to a file system hosted by a secure EVS. Similarly, you cannot configure an individual security context for an EVS (turning an EVS into a secure EVS) if there are CNS links to one or more file systems in that EVS.
 - **Name space usage and the secure EVS.** If you want to use a name space with a secure EVS that does not use the global configuration settings, you must configure an EVS name space for that secure EVS. An EVS name space is required because file systems hosted by the secure EVS cannot be linked to from the CNS, and file systems hosted by the EVS cannot access links in the CNS.
If you want to use a name space with a secure EVS that does use the global configuration settings, you may configure an EVS name space for that secure EVS, but it is not required. If the secure EVS does use the global security settings, the file systems hosted by the EVS can access links in the CNS.

- **Links to a secure EVS individual name space.** In a secure EVS with an individual name space, you can add links between file systems hosted by the same secure EVS.
- **Configuring a group of EVSs with the same settings.** To create a group of secure EVSs that use the same individual security context settings (that are different from the global settings), you must configure each secure EVS in the group separately.
- **Reconfiguring a secure EVS security context to use the global context.** If a secure EVS is reconfigured to use the global security context (reverting it to an EVS), and the secure EVS was using a different NT domain than the cluster, CIFS names and CIFS share names become invalid. This occurs because CIFS names (and CIFS share names) are associated with a specific NT domain, and the NT domain name changes. If the global security context and the secure EVS and the NT domain are different, after you remove the secure EVS' individual security context (making it an EVS again), you must delete all CIFS names for the EVS and all CIFS shares for the file systems in the EVS. Then, you must recreate the EVS CIFS names and the CIFS shares for the file systems in the EVS.

About security contexts

Because EVSs and secure EVSs inherit many of their settings from the cluster's global context, when configuring name services, you must specify if you want to change the global context or the individual context.

For example, on the **NIS/LDAP Configuration** page or the **EVS Details** page, the current security context displayed as follows.

EVS Security Context: Global Configuration

change...

The EVS security context can be any of the following:

- **Global Configuration**, which indicates that the current security context is the global context.
- **Inherits Global Configuration**, which indicates that the current security context is an individual EVS security context that has been set to use the settings that are defined in the global security context.
- **Individual Configuration**, which indicates that the current security context is an individual context that has individually specified settings.

On the **Name Services** and **NIS/LDAP Configuration** pages, click change to switch the current context between an individual context for a particular secure EVS and the global context.

If you make changes that affect:

- The global context, those changes apply to all EVSs that have security context settings set to Inherits Global Configuration.

- An individual context, those changes apply only to the currently selected EVS.

On the **EVS Details** page, click change to switch the context used by the EVS:

- If an individual context is being used, you can change the EVS to use the global context, removing the individual security context and changing the secure EVS into a regular EVS.
- If the global context is being used, you can change the EVS to use an individual context (creating a secure EVS).

Security context contents

The following security settings make up the security context for all EVSs and secure EVSs:

- Name services (DNS, WINS, and so on)
- Windows NT domain
- NIS domain
- User/group/domain mapping tables
- Security mode (mixed or UNIX)
- Individual groups
- **cifs-auth** setting
- **bypass-permissions-checks** setting
- File/directory umasks
- NFS export and CIFS share access options
- CNS mount points

The parts of the security context that can be configured for a secure EVS include:

- Name services (DNS, WINS, and so on)
- Windows NT domain
- CNS mount points

Securing an EVS

To change an EVS into a secure EVS, you must add an individual security context.

Procedure

1. Ensure the EVS Security license key is installed.
Before you can specify an individual security context for an EVS, the secure EVS license must be installed.

2. Navigate to **Home > Server Settings > EVS Management** to display the **EVS Management** page.

Server Settings [Home](#) > [Server Settings](#) > [EVS Management](#)

EVS Management

Filter

No Filtering Applied

[filter](#)

▼ Label	Type	Cluster Node	Status	First IP Address	First Port	
<input type="checkbox"/> donotdelete	File Services	Group1-node2	Online	172.31.60.45/24	ag1	details
<input checked="" type="checkbox"/> EVS1	File Services	Group1-node1	Online	172.31.60.47/24 ...	ag1 ...	details
<input type="checkbox"/> g1-eva1	File Services	Group1-node1	Online	172.31.60.41/24	ag1	details
<input type="checkbox"/> g1-eva2	File Services	Group1-node1	Online	172.31.60.42/24	ag1	details
<input type="checkbox"/> g1-eva3	File Services	Group1-node1	Online	172.31.60.43/24	ag1	details
<input type="checkbox"/> Group1-admin	admin services	Group1-node1	Online	192.0.2.3/24 ...	eth1 ...	details
<input type="checkbox"/> LNAS	File Services	Group1-node1	Online	172.31.60.48/24	ag1	details

[Check All](#) | [Clear All](#)

Actions: [enable](#) [disable](#) | [add](#)

Shortcuts: [IP Addresses](#) [EVS Migration](#) [Link Aggregation](#)

3. Fill the check box for the EVS you want to disable, and click **disable**.

- Click the **details** button next to EVS you want to change into a secure EVS.

Server Settings [Home](#) > [Server Settings](#) > [EVS Management](#) > EVS Details

EVS Details EVS1

Name: [rename](#)

EVS ID: 6

Status: Online

Type: File Services

Enabled: Yes

Preferred Cluster Node: [apply](#)

EVS Security: Global [change...](#) (Disable EVS to alter EVS security)

Default File System Security Mode: [Mixed \(Windows and Unix\)](#)

File Systems

[FS11](#)

IP Addresses

Port	IP Address
ag1	172.31.60.47/24
ag1	face::17/64

Actions: [enable](#) [disable](#) | [delete](#)

Shortcuts: [IP Addresses](#) [EVS Migration](#)

- Click **change** to add an individual security context to the EVS.
- Click **OK** to confirm the change, or **cancel** to return to the **EVS Details** page.
- Select the EVS, and click **enable**.
- Recreate CIFS names for the secure EVS.
For information on specifying a CIFS name for the secure EVS, refer to the *File Services Administration Guide*.
- Adjust CIFS shares for the file systems in the secure EVS.
For information on recreating CIFS shares for the file systems in the secure EVS, refer to the *File Services Administration Guide*.
- Specify user and group access for the secure EVS.
For information on configuring user and group access to file systems in the secure EVS, refer to the *File Services Administration Guide*.
- Configure name services for the secure EVS.
For information on configuring name services for the secure EVS, refer to the *Network Administration Guide*.
- Configure the EVS name space, if necessary.
If you want to use a name space with a secure EVS that does not use the global configuration settings, you must configure an EVS name space for that secure EVS.

Removing an individual security context from a secure EVS

Removing an individual security context from an EVS changes the secure EVS back into an EVS.

Procedure

1. Navigate to **Home > Server Settings > EVS Management** to display the **EVS Management** page.

EVS Management

Filter: No Filtering Applied

Label	Type	Cluster Node	Status	First IP Address	First Port	
<input type="checkbox"/> donotdelete	File Services	Group1-node2	Online	172.31.60.45/24	ag1	details
<input checked="" type="checkbox"/> EVS1	File Services	Group1-node1	Online	172.31.60.47/24 ...	ag1 ...	details
<input type="checkbox"/> g1-evs1	File Services	Group1-node1	Online	172.31.60.41/24	ag1	details
<input type="checkbox"/> g1-evs2	File Services	Group1-node1	Online	172.31.60.42/24	ag1	details
<input type="checkbox"/> g1-evs3	File Services	Group1-node1	Online	172.31.60.43/24	ag1	details
<input type="checkbox"/> Group1-admin	admin services	Group1-node1	Online	192.0.2.3/24 ...	eth1 ...	details
<input type="checkbox"/> LNAS	File Services	Group1-node1	Online	172.31.60.48/24	ag1	details

[Check All](#) | [Clear All](#)

Actions: [enable](#) [disable](#) | [add](#)

Shortcuts: [IP Addresses](#) [EVS Migration](#) [Link Aggregation](#)

2. Fill the check box for the secure EVS you want to disable, and click **disable**.

3. Click the **details** button next to secure EVS you want to change into an EVS.

Server Settings [Home](#) > [Server Settings](#) > [EVS Management](#) > EVS Details

EVS Details EVS1

Name: [rename](#)

EVS ID: 6

Status: ● Online

Type: File Services

Enabled: Yes

Preferred Cluster Node: [apply](#)

EVS Security: Global [change...](#) (Disable EVS to alter EVS security)

Default File System Security Mode: [Mixed \(Windows and Unix\)](#)

File Systems

[FS11](#)

IP Addresses

Port	IP Address
ag1	172.31.60.47/24
ag1	face::17/64

Actions: [enable](#) [disable](#) | [delete](#)

Shortcuts: [IP Addresses](#) [EVS Migration](#)

4. Click **change** to remove an individual security context from the secure EVS.
5. Click **OK** to confirm the change, or **cancel** to return to the **EVS Details** page.
6. Select the EVS, and click **enable**.

EVS name spaces

An EVS name space allows separate file systems within a virtual server (EVS) to appear as subdirectories of a single logical file system (that is, as one unified file system). An EVS name space can also make multiple storage elements on the virtual server available to network clients through a single CIFS share or NFS export.

The EVS name space functions in the same way as the cluster name space (CNS), except that its context is that of the EVS, instead of the cluster.

In order to create an EVS name space, you must have installed a CNS license, and an EVS Security license, and you must have set the EVS to use an individual security context.

Linking to and from an EVS name space has the following constraints:

- **Links within an EVS name space.** In an EVS name space tree, you can add links from the EVS name space to file systems hosted by the same secure EVS.
- **Links between the CNS and the EVS name spaces.** The contexts of the Cluster Name Space and the EVS name space are mutually exclusive: links from one to the other are not allowed.
- **Links outside the EVS name space.** Links from the individual EVS name space to file systems in other EVSs are not supported.

Creating an EVS

Before they can be used, EVSs must be created and assigned to an IP address; then, to provide file services, assign one or more file systems.

Procedure

1. Navigate to **Home > Server Settings > EVS Management**, and click **Add EVS** to display the **Add EVS** page.

2. Enter the requested information (all fields are required).
3. Click **OK** to save, or **cancel** to decline.

Assigning a file system to an EVS

After the EVS has been created, at least one file system must be assigned to it. You can either create a new file system on the EVS (see the *File Services Administration Guide*), or you can assign an existing file system to the EVS. To assign a file system to an EVS, you can relocate a file system currently assigned to another EVS (see the *Replication and Disaster Recovery Administration Guide*) or you can assign a file system that is currently not assigned to an EVS.

Procedure

1. Navigate to **Home > Storage Management > File Systems** to display the **File Systems** page.

2. In the file system grid, for the specific file system that will be assigned to the EVS, click **details** to display the **File System Details** page.

File System Details

Storage Management [Home](#) > [Storage Management](#) > [File Systems](#) > File System Details

Label:

Capacity

53% Total Used Space

Capacity: 489.31 GB
Free: 231.10 GB (47%)
Total Used: 258.21 GB (53%)
Expansion Limit: 500 GB
Auto-expansion is disabled

Legend: ☒ Live file system ☐ Usage Warning ☐ Usage Severe

Configuration

Status: Mounted
Deduplication: [Disabled](#)
Thin Provisioning: Disabled
EVS: g1-evs3 (Online)
Security Mode: [Mixed \(Windows and Unix\) \(Inherited\)](#)
Block Size: 32 KB
Read Cache: No
WFS Version: WFS-2
Syslock: Disabled
Object Replication Target: Disabled
Transfer Access Points During Object Replication: Enabled

Usage Thresholds

File System Usage

	Live file system	Snapshots	Entire file system
Current:	53 %	0 %	53 %
Warning:	90 %	90 %	95 %
Severe:	97 %	97 %	97 %

☐ Do not allow the [live](#) file system to expand above its Severe limit

Associations

Storage Pool: [SP_1](#)

Capacity: 42.72 TB
Free: 29.77 TB (70 %)
Used: 12.95 TB (30 %)

Related File Systems

Data Migration From: g1-cluster PHDS1
File Replication To: gizmo1 SiteACopy/
Last: 2014-06-03 00:00
Last: Not run

Check / Fix

Status: File System is not being checked or fixed.
Scope: ☐ Entire file system ☒ Directory Tree
 [Active Tasks](#)

Actions:

Shortcuts: [Data Migration Paths](#) [File System Versions](#)

3. From the EVS list, select the EVS to which you want to assign the file system, and click **assign**.
4. When the **Server Settings** page is displayed, in the file system list, verify that the new assignment is displayed in the EVS column; select the file system, and click **mount**.
The file system status changes to *Mounted*.

Virtual server (EVS) management

The EVS Management page allows EVSs to be added, enabled, and disabled.

Procedure

1. Navigate to **Home > Server Settings > EVS Management** to display the **EVS Management** page.

[Server Settings](#) [Home](#) > [Server Settings](#) > [EVS Management](#)

EVS Management

Filter

No Filtering Applied


[filter](#)

▼ <u>Label</u>	<u>Type</u>	<u>Cluster Node</u>	<u>Status</u>	<u>First IP Address</u>	<u>First Port</u>	
<input type="checkbox"/> donotdelete	File Services	Group1-node2	Online	172.31.60.45/24	ag1	details
<input checked="" type="checkbox"/> EVS1	File Services	Group1-node1	Online	172.31.60.47/24 ...	ag1 ...	details
<input type="checkbox"/> g1-eva1	File Services	Group1-node1	Online	172.31.60.41/24	ag1	details
<input type="checkbox"/> g1-eva2	File Services	Group1-node1	Online	172.31.60.42/24	ag1	details
<input type="checkbox"/> g1-eva3	File Services	Group1-node1	Online	172.31.60.43/24	ag1	details
<input type="checkbox"/> Group1-admin	admin services	Group1-node1	Online	192.0.2.3/24 ...	eth1 ...	details
<input type="checkbox"/> LNAS	File Services	Group1-node1	Online	172.31.60.48/24	ag1	details

[Check All](#) | [Clear All](#)

Actions: [enable](#) [disable](#) | [add](#)

Shortcuts: [IP Addresses](#) [EVS Migration](#) [Link Aggregation](#)

Field/Item	Description
Filter	Allows you to filter the list of EVSs by EVS name, status, or cluster node.
Label	EVS name (identifier).
Type	Type of service: administrative services or file services.
Cluster Node	Cluster node on which the EVS is currently residing (only displayed for cluster nodes).
Status	Service status: <ul style="list-style-type: none"> • Online: Up and accessible. • Disabled: Down and inaccessible.
IP Address	First IP address assigned to the EVS. <div>  Note: An EVS can have multiple IP addresses. </div>
Port	The Ethernet port or aggregation to which the IP address for the EVS is assigned.
details	Displays the EVS Details page for the selected EVS.

Field/Item	Description
Check All	Selects all EVSs in the list above.
Clear All	Deselects all EVSs in the list above.
Actions	
enable	Enables a disabled EVS.
disable	Disables and enabled EVS.
add	Creates a new EVS.
Shortcuts	
IP Addresses	Opens the IP Addresses page.
EVS Migration	Opens the EVS Migration page.

Displaying EVS details

Procedure


1. Navigate to **Home > Server Settings > EVS Management**, and click **details** to display the **EVS Details** page.

Server Settings [Home](#) > [Server Settings](#) > [EVS Management](#) > EVS Details

EVS Details EVS1

Name: [rename](#)

EVS ID: 6

Status:  Online

Type: File Services

Enabled: Yes

Preferred Cluster Node: [apply](#)

EVS Security: Global [change...](#) (Disable EVS to alter EVS security)

Default File System Security Mode: [Mixed \(Windows and Unix\)](#)

File Systems


[FS11](#)



IP Addresses


Port	IP Address
ag1	172.31.60.47/24
ag1	face::17/64

Actions: [enable](#) [disable](#) | [delete](#)

Shortcuts: [IP Addresses](#) [EVS Migration](#)

Field/Item	Description
Name	EVS identifier (same as Label in the previous page).
EVS ID	Unique identifier for the EVS within the cluster, generated by the server upon EVS creation. <div> Note: If moved to another server in a server farm, the EVS ID might change, but not if the move is within a cluster.</div>
Status	Service status: <ul style="list-style-type: none">• Online: Up and capable of providing services.• Offline: Not running. While offline, EVS are inaccessible.

Field/Item	Description
Type	Type of service provided by the EVS: administration services or file services.
Enabled	Yes (enabled), or No (disabled).
Preferred Cluster Node	Preferred cluster node for the EVS (only displayed for file serving EVSs in a cluster.) Indicates cluster node preference; the EVS might, however, be hosted by a node other than its preferred node after having been migrated, for several reasons, such as node failure, manual migration for load balancing, and so forth.
EVS Security	<p>Displays the current EVS security context. Click change to select a different EVS security context, or to select the global configuration.</p> <p>Selecting a different EVS security context changes how EVS name services, user mappings, group mappings, localgroups, name space (CNS), DNS servers, and NIS/LDAP configuration settings are managed.</p> <ul style="list-style-type: none"> • If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS. • If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To change the settings of an EVS using an individual security context, you must go to the configuration settings for the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context. <hr/> <p> Note: You must disable the EVS before you can switch the EVS security context being used (between the global security context and an individual security context).</p> <hr/>
Default File System Security Mode	Displays the default security model to be used by file systems in the EVS.
File Systems	<p>List of all file systems hosted by the EVS. For more information about a particular file system, click its name to display its File System Details page.</p> <hr/> <p> Note: Administrative EVSs (EVS Type set to Admin Services), the this area is not displayed</p> <hr/>
IP Addresses	The IP addresses area displays a list of all IP addresses assigned to the EVS. Note that an EVS can have multiple IP addresses.

Field/Item	Description
Subnet Mask	Subnet mask for the EVS.
Port	The cluster node gigabit Ethernet port to which the IP address for the EVS is assigned.
Actions	
rename	Applies the new label entered in the Name field.
apply	Applies a new preferred cluster node selected from the Preferred Cluster Node list.
enable/disable	Controls the status of the EVS.
delete	Removes the EVS. Do not click delete until you first disable the EVS. <div>  Note: Deleting an EVS does not affect the file system owned by the EVS. After the EVS has been deleted, assign the file system to another EVS to make it available for use. </div>

Migrating an EVS within a cluster

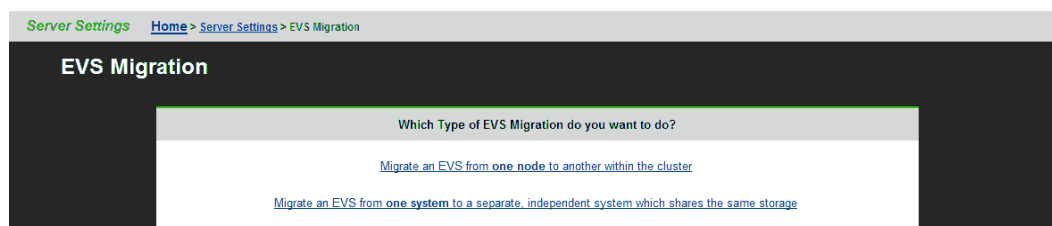
While migration of EVSs occurs automatically as part of the failover resiliency of a cluster, EVSs can be manually migrated to a different node in a cluster, or among servers or clusters within a server farm.

An individual EVS can be migrated to a different node within the same cluster, or all EVSs can be migrated to another server or another cluster. The current mapping of EVSs to cluster nodes can be preserved, and the saved map is called a preferred mapping.

Procedure

1. Navigate to **Home > Server Settings > EVS Migration** to display the **EVS Migration** page.

If the currently managed server is in a cluster and the SMU is also managing at least one standalone server, the following page displays (if the SMU is not managing a cluster and one or more standalone servers, this page does not display):



2. Select **Migrate an EVS from one node to another within the cluster** to display the main **EVS Migration** page.

Server Settings [Home](#) > [Server Settings](#) > EVS Migration

EVS Migration

EVS Mappings

Node	Current EVS Mapping	Preferred EVS Mapping
Group1-node1	g1-evs3 , g1-evs1 , LNAS , g1-evs2 , EVS1	g1-evs3 , LNAS , g1-evs2 , EVS1
Group1-node2	donotdelete	donotdelete , g1-evs1

[Save current](#) as preferred | [Migrate all](#) to preferred

An orange EVS indicates the EVS is **not** on its preferred cluster node.


A grey EVS indicates the EVS does **not** have a preferred cluster node.

A black EVS indicates the EVS is on its preferred cluster node.

EVS Migrations

☐ Migrate EVS Group1-admin to cluster node Group1-node2

☐ Migrate all EVSes from cluster node Group1-node1 to cluster node Group1-node2

 Migrating the EVS will disrupt file system services to any existing clients.

migrate

Shortcuts: [File System Ops/Sec](#) [EVS Management](#) [Cluster Configuration](#)

3. Perform a migration of the type required.
 - To migrate all EVSs between cluster nodes:
 1. Select **Migrate all EVS from cluster node ____ to cluster node ____**.
 2. From the first list, select the cluster node from which to migrate all EVS.
 3. From the second list, select the cluster node to which to migrate all EVS.
 4. Click **Migrate**.
 - To migrate an EVS to a cluster node:
 1. Select **Migrate EVS ____ to cluster node ____**.
 2. From the first list, select the cluster node to migrate.
 3. From the second list, select the cluster node to which the EVS will be migrated.
 4. Click **Migrate**.
 - To save a preferred EVS to cluster node mapping:



Note: Saving the current EVS-to-cluster configuration as the preferred mapping helps when restoring EVSs to cluster nodes. For example, if a

failed cluster node is being restored, the preferred mapping can be used to easily restore the original cluster configuration.

1. Migrate the EVS between the cluster nodes until the preferred mapping has been defined. The current mapping will be displayed in the Current EVS list box.
 2. To save current EVS-to-cluster node mapping, click **Save** in **Save current EVS mapping as the preferred mapping**. The preferred mapping will then be displayed in the Current EVS column.
- To migrate all EVSs to a preferred mapping, click **Migrate** in **Migrate all EVS to their preferred mapping**.

Migrating an EVS within a server farm

Migration within a server farm is supported under the following conditions:

- When both the source and destination server are online.
- If the source server is offline and the destination server is online.
- The EVS does not contain any file systems that are linked into a CNS tree.
- When both the source and destination server have the same major firmware revision.

After migrating EVS between servers in a server farm, the assignment of tape drives and tape autochanger devices to EVS must be manually adjusted:

- Tape devices specifically assigned to a migrated EVS will have become unassigned.
- Tape devices assigned to "any EVS" on the source server will remain assigned to "any EVS" on the source server.

Tape devices must not be assigned to EVSs on more than one server.

While EVSs contain most of the settings required to support client storage access, some settings (including DNS, Windows NT domain and Active Directory) are functions of the host server or node, not of the EVS itself. Therefore, when preparing to migrate an EVS from one server or cluster to another, verify in advance that the target server's settings can properly support the EVS. To prepare a server or cluster to receive a new EVS, source server settings can be cloned.

Cloning server settings

Procedure

1. Navigate to **Home > Server Settings**, select the target of the migration as the SMU's managed server, and click **Clone Server Settings** to display the **Clone Server Settings** page.

2. From the list, select the server/node currently hosting the EVS, and click **next**.
3. Select the settings to clone to the target server/node.



Caution: Settings selected for cloning will overwrite the currently defined settings on the target server/node. To keep part of the existing configuration, do not clone the configuration items you want to keep.

4. Click **OK** to initiate cloning.
After the target server/node has been prepared through server cloning, it is ready to receive the EVS migration.

Migrating an EVS within a server farm

Procedure

1. Navigate to **Home > Server Settings > EVS Migration**.

This page will only display if the currently managed server is a cluster node. Otherwise, clicking **EVS Migrate** will immediately launch the **EVS Migration** page.

2. Click **Migrate an EVS from one system to a separate, independent system** to display the **EVS Migration** page.
3. Click **change** to select a source server and source EVS.
4. From the **Destination Server** menu, select a target server.
5. Fill the **Test Only** check box to test the migration before committing the change.
This ensures that the EVS migration is possible.



Note: When selecting a destination server for an EVS, note that both the source and destination server must be running the same major firmware revision.

A message displays indicating if the operation succeeded or failed. If the operation failed, the message includes the reason for the failure.

6. Click **Migrate** to start the process.



Note: If the source server is offline or doesn't function, migration will be performed using an existing backup and a warning is displayed.

Status and monitoring

Web Manager provides comprehensive and integrated management of the storage server and its storage subsystem. Its management pages provide color-coded information about the status of the various installed devices. Web Manager also provides a comprehensive event logging and alerting mechanism, which can notify the system administrator, as well as Hitachi Data Systems Support Center, as soon as a problem occurs. Alerts are issued through email, SNMP, syslog, and Windows pop-ups.

The SMU can also publish information about system drives (SDs) on HDS storage subsystems to the Hitachi Device Manager (HDvM). This integration allows the Hitachi Tiered Storage Manager (HTSM) to retrieve information about storage devices attached to a Hitachi NAS Platform or cluster by searching for that information on the HDvM.

File system auditing monitors and records file system operations performed through the CIFS protocol. File system operations such as file access and deletion are recorded in the server's file system audit log. You can then display the file system audit log through a remote Windows Event Viewer, and save the log entries in `.evt` format for later review.

- ☐ [Storage system status](#)
- ☐ [Configuring devices on the System Monitor](#)
- ☐ [Checking the system status](#)
- ☐ [Using the server status console](#)
- ☐ [Checking the status of a server unit](#)
- ☐ [Checking UPS status](#)
- ☐ [Checking SMU status](#)
- ☐ [Monitoring multiple servers](#)

- [Monitoring storage subsystems with Hitachi Device Manager](#)

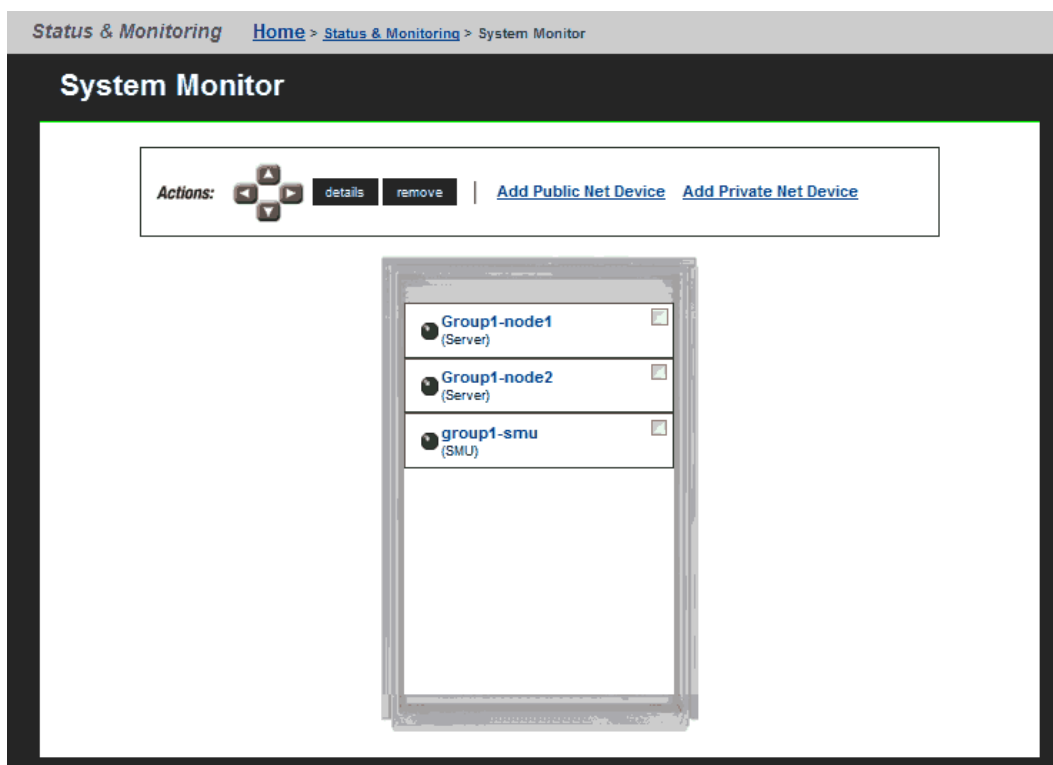
Storage system status

Web Manager provides a flexible, customizable, and easy-to-use interface, displaying the status of each managed device in the storage system. Ethernet-connected auxiliary devices can be added to the System Monitor as managed objects, so that the status of these devices will be displayed. The **System Monitor** page provides a central management console for the management and status monitoring of all devices that comprise the network storage system.

Configuring devices on the System Monitor

Procedure

1. Navigate to **Home > System Monitor** to display the **System Monitor** page.



2. Optionally, change the position of any components by filling its check box, and using the arrows in the Action section.
3. Optionally, display the status or details for any component.
The rows in the following table list the basic components that make up a Hitachi NAS Platform system. This table indicates what happens when you click a component's name in the component list.

4. The following Actions are available and apply to selected components:
- Click **remove** to delete a component.
 - Click **details** to display details regarding a particular component.
 - Click add **Public Net Device** to add a device residing on the public (data) network.
 - Click add **Private Net Device** to add a device residing on the private (management) network.



Note: Devices on the private management network are “hidden” from the data network through Network Address Translation (NAT).

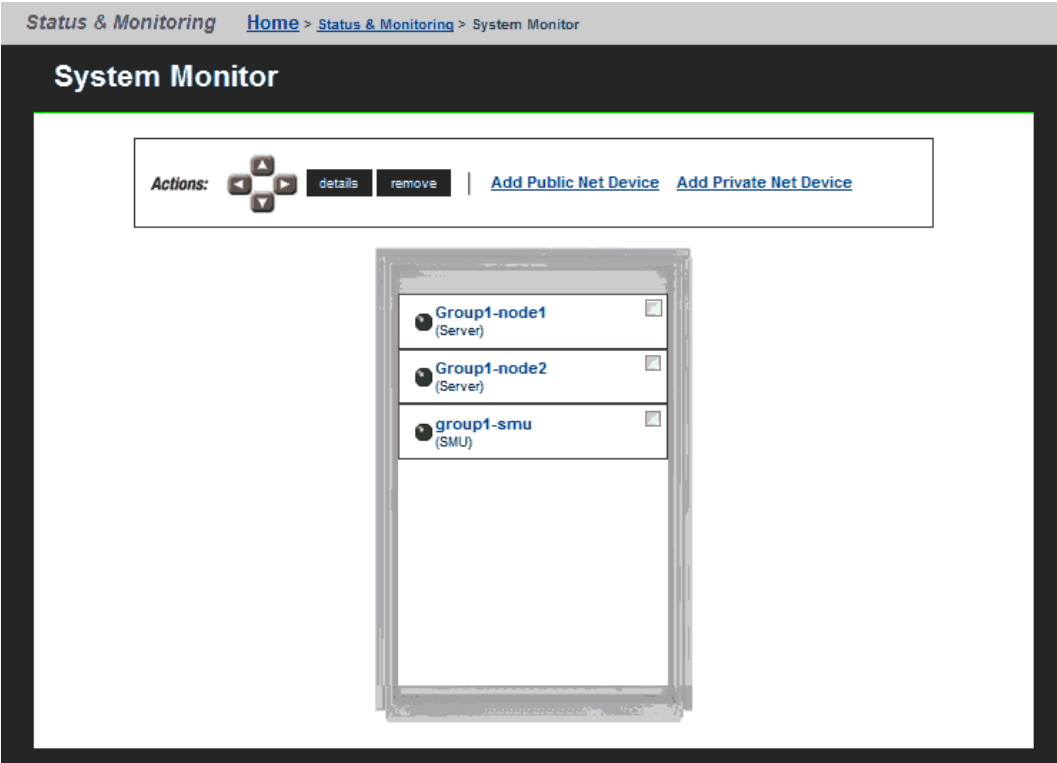
After a device is added to the System Monitor, clicking its name:


- Opens its embedded management utility in the Web browser, using either HTTP, HTTPS, or Telnet.
- The SMU periodically checks for device activity and connectivity with the server; if a device fails to respond to network “pings”, the System Monitor changes its color to red and the SMU issues an alert (devices can also be configured to send SNMP traps to the SMU).
- Events from the device will be added to the event log if the SMU has a MIB for the device.

Checking the system status



Procedure



1. Navigate to **Home > Status & Monitoring > System Monitor** to display the **System Monitor** page.



 **Note:** The System Monitor reflects a 60-second delay for status information cached by the SMU.

When displaying a device's status using the colored LED, the following conventions apply:

Color	Status	Description
 Green	Information	Operating normally and not displaying an alert condition.
 Amber	Warning	Needs attention, but does not necessarily represent an immediate threat to the operation of the system.

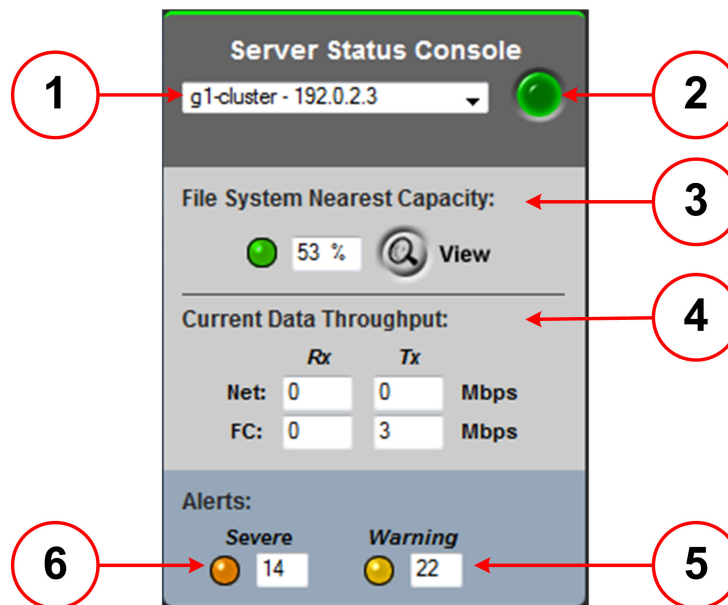
Color	Status	Description
 Red	Critical	Requires immediate attention. The failure is critically affecting the operation of the system.
 Gray	Unknown	Status of the device cannot be determined; for example, if the server is out of contact with the SMU, the server and the status of its components cannot be determined.

Using the server status console

Summary status information for the currently managed server can be displayed from the Web Manager's **Server Status Console**.

Procedure

1. Navigate to the **Home** page to locate the **Server Status Console**, which displays summary status information pertaining to the currently managed server.



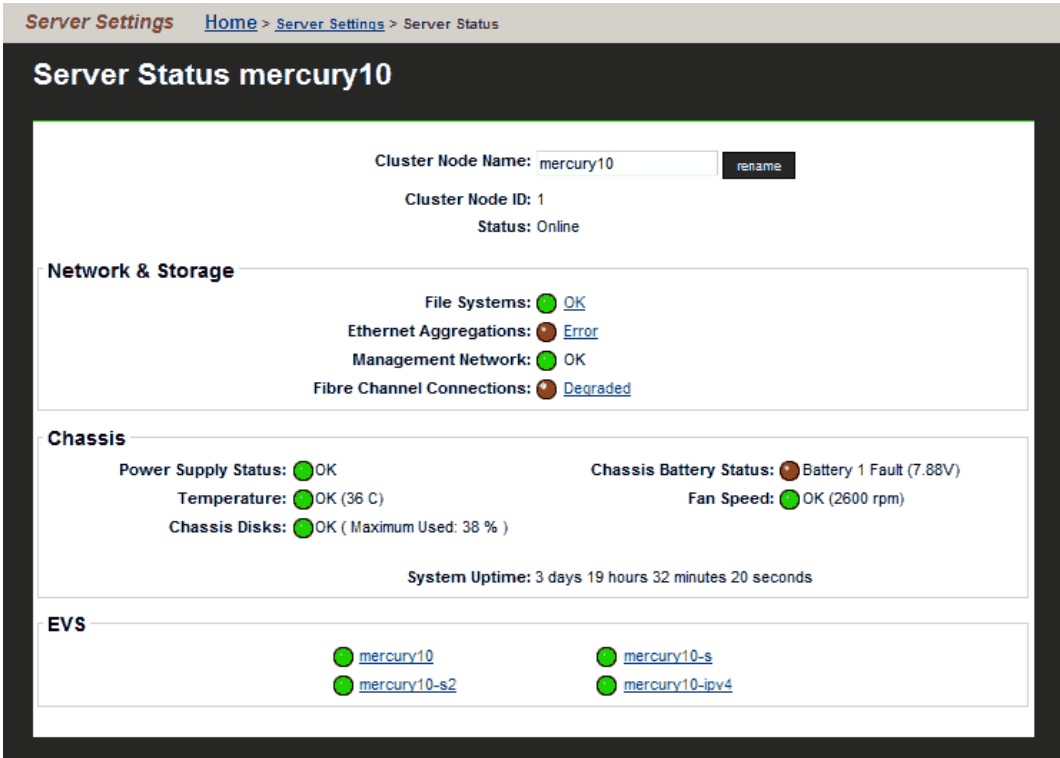
Item	Description
1	Currently managed server or cluster name and IP address. Use the drop-down list (of managed servers/clusters) to select another server or cluster to manage.
2	Summary status indicator for the currently managed server or cluster: <ul style="list-style-type: none"> • <i>Green</i>—Operating normally (not showing an alert condition). • <i>Amber</i>—Warning condition (operating normally, but action should be taken to maintain normal operation). • <i>Red</i>—Critical condition (not functioning or failing in a way that presents a danger to the system).
3	File System Nearest Capacity The color of the status indicator color provides information about how close the file system is to its maximum configured size limit. <ul style="list-style-type: none"> • <i>Green</i>—Usage is below the Warning threshold. • <i>Yellow</i>—Usage has reached or exceeded the Warning threshold, but is below the Critical threshold. • <i>Orange</i>—Usage has reached or exceeded the Critical threshold. The percentage of allocated space used by the file system nearest to full capacity is displayed next to the status indicator. Click View to display the File Systems page, where you can find out more about the file systems on the server or cluster.

Item	Description
4	<p>Current Data Throughput</p> <p>Displays the current data throughput from the data network and Fibre Channel ports, both received (Rx) and transmitted (Tx). Throughput values are updated every 10 seconds.</p>
5	<p>Warning events recorded in the event log during the past 24 hours.</p> <p>Click the yellow indicator to display the event log, showing all warning events recorded in the event log.</p>
6	<p>Severe events recorded in the event log during the past 24 hours.</p> <p>Click the orange indicator to display the event log, showing all severe events recorded in the event log.</p>

Checking the status of a server unit

Procedure

1. Navigate to **Home > Status & Monitoring > Server Status**.
 - For a stand-alone server that is not part of a cluster, the **Server Status** page is displayed.



Field/Item	Description
Cluster Node Name	The server/node name (label). Click rename and enter a new name, to change the server name.
Cluster Node ID	The ID assigned to the node.
Status	Indicates the node status: <ul style="list-style-type: none">• Online: Node has completed booting.• Unknown: Node has not yet booted.• Up: Node is booting (displayed only while the node is booting).• Not up: node is shutting down (displayed only while the node is shutting down).• Dead: Node has failed to go online after booting
Network and Storage	

Field/Item	Description
File Systems	<p>Overall indicator of file system status:</p> <ul style="list-style-type: none"> • OK. All file systems up and operational. • Failed. One or more file systems has failed. <p>Click the status link to display the File Systems page, which lists all the file systems assigned to the EVS in that cluster node.</p>
Ethernet Aggregations	<p>Overall status of Ethernet aggregations in the server/node:</p> <ul style="list-style-type: none"> • OK. All aggregated ports are up and linked. • Degraded. One or more ports in an aggregation has failed. • Failed. All ports in an aggregation have failed. <p>Click the status link to display the Link Aggregation page, which lists all aggregations (trunks) in the server/node.</p>
Management Network	<p>Overall status of the management network:</p> <ul style="list-style-type: none"> • OK. Links are up and heartbeats are being received. • Failed. No heartbeats are being received, and the links may be up or down. <p>Click the status link to display the Ethernet Statistics page, which lists information about the management port and the aggregated Ethernet ports in the server/node.</p>
Fibre Channel Connections	<p>An overall status indicator for the Fibre Channel ports in the server/node:</p> <ul style="list-style-type: none"> • OK. All ports up and operational. • Degraded. Some ports up and operational, but one or more has failed. • Failed. All ports have failed. <p>Click the status link to display the Fibre Channel Statistics Per Port page, which lists all Fibre Channel ports in use in the server/node.</p>
Chassis	
Power Supply Status	<p>A status indicator for the power supply units (PSUs):</p> <ul style="list-style-type: none"> • OK. Both PSUs are installed and operating normally. • Not Fitted. One PSU not responding to queries, which may mean that has been removed from the chassis, or is not properly installed in the chassis. • Fault or Switched Off. One PSU not responding to queries, and it has failed, been switched off, or is not plugged in to mains power. • Unknown. One PSU not responding to queries, and the exact cause cannot be determined.
Temperature	<p>Status indicator for temperature of the server/node chassis:</p> <ul style="list-style-type: none"> • OK. Within the normal operating range. • Degraded. Above normal, but not yet critical. • Failed. Critical. <p>When available, the temperature in the chassis also is displayed. The displayed temperature is the highest reported temperature of any of the boards in the chassis.</p>

Field/Item	Description
Power Supply Battery Status	<p>Status of the power supply battery.</p> <p>When the indicator is green:</p> <ul style="list-style-type: none"> • OK. Capacity and voltage within the normal operating range. • Initialising. PSU battery is initializing after initial installation. • Normal Charging. PSU battery is being charged. • Cell-Testing. PSU battery is being tested. <p>When the indicator is amber:</p> <ul style="list-style-type: none"> • Discharged. Capacity and/or voltage below normal. This status should be considered a warning; if it continues, the PSU battery should be replaced. • Low. Capacity or voltage below normal operating level. This status should be considered a warning; if it continues, the PSU battery should be replaced. • Not Responding. PSU battery is not responding to queries. <p>When the indicator is red:</p> <ul style="list-style-type: none"> • Fault. PSU battery is not holding a charge, has the wrong voltage, or some other fault, and the PSU battery should be replaced. • Not Fitted. PSU battery is not detected. Contact your technical support representative for more information. • Failed. Capacity and voltage consistently below acceptable minimum, or the PSU battery is not charging, or is not responding to queries. This status indicates a failure; the PSU battery should be replaced. • Very Low. Capacity and voltage below acceptable minimum. If this status continues for more than a few hours, it indicates a failure; the PSU battery should be replaced. <p>When available, the level of the battery charge also is displayed.</p>
Fan Speed	<p>Status of fans in the server/node chassis:</p> <ul style="list-style-type: none"> • OK. All fans operating normally. • Degraded. One or more fans spinning below normal range. • Failed. At least one fan has stopped completely, or is not reporting status. <p>When available, the chassis fan speed also is displayed. The displayed fan speed is the slowest reported speed of any of the three fans. An error message might be displayed, even if it does not correspond with the slowest fan.</p>
System Uptime	Duration since last reboot of the server/node.
EVS	
EVS	<p>Displays the names (labels) of EVSs assigned to the node, and displays a status indicator for each EVS:</p> <ul style="list-style-type: none"> • Green. Online and operational. • Amber. Offline, but listed here because the server/node is hosting the administrative EVS. • Red. Failed. <p>Click the EVS name to display the EVS Details page for that EVS.</p>

- For a cluster node, the **Cluster Configuration** page is displayed.

Server Settings [Home](#) > [Server Settings](#) > Cluster Configuration

Cluster Configuration

Cluster Nodes						EVS	
Name	IP Address	Status	Model	Health			
Group1-node1	192.0.2.200	Online	3090-G2	OK			
Group1-node2	192.0.2.201	Online	3090-G2	OK		Group1-admin, g1-evs2, donotdelete, g1-evs1, LNAS, g1-evs2, EVS1	details

Cluster Information		Quorum Device	
Cluster Name: g1-cluster	rename	Name: GROUP1-SMU	
Status: Online		IP Address: 192.0.2.1	
Health: Robust		Status: Configured	
Cluster UUID: a6e6ddf0-9627-11cb-9000-d428dd593ca4		add remove	
MAC: d4-28-d5-99-3c-a4			

Actions: [Add Cluster Node](#)

Shortcuts: [Quorum Services v2](#) [EVS Management](#) [EVS Migration](#)

Field/Item	Description
Cluster Nodes	
Name	Node name.
IP Address	IP address of the cluster node. This IP address is on the private management network, which connects devices within the cluster.
Model	Server model, if available.
Health	<p>Worst-case status from each node:</p> <ul style="list-style-type: none"> OK. Operating normally, with no failures. Degraded. Operating, but with one or more failures in connectivity. The problem might be with the cluster interconnect, the management network, or quorum device communication. Failed. Not operating, due to one or more failures in connectivity. The problem might be with the cluster interconnect, the management network, or the quorum device communication. <p>This page also shows the status of operations of the server's internal hard disks, and the percentage of the server's internal disk space that has been used. Disk status is shown as:</p> <ul style="list-style-type: none"> OK. Operating normally. Degraded. A non-critical problem has been discovered with one or both of the server's internal hard disks. Failed. A critical problem has been discovered with one or both of the server's internal hard disks.
EVS	Displays the names (labels) of EVSs hosted on each cluster node. Click the EVS name to display the EVS Details page for that EVS.
Cluster Information	
Cluster Name	Cluster name. Click rename , and enter a new name in the field, to rename the cluster.
Health	<p>Cluster health:</p> <ul style="list-style-type: none"> Robust. Operating normally, with no failures in the cluster interconnect, the management network, or quorum device communication. Degraded. Operating, but one or more nodes has failed or there has been a failure in connectivity.

Field/Item	Description
Cluster UUID	UUID (unique ID) of the cluster. This string provides a unique identifier for each cluster when there are several clusters on a network.
MAC	MAC address of the cluster.
Quorum Device	
Name	SMU on which the QD resides.
IP Address	IP address of the SMU on which the QD resides.
Status	QD status: <ul style="list-style-type: none"> • Configured. Attached to the cluster, but vote not needed. The QDs vote is not needed when any cluster contains an odd number of operational nodes. • Owned. Attached to the cluster and owned by a specific cluster node. • Not up. Cannot be contacted. • Seized. Taken over by another cluster.

Checking UPS status

Displays the status of a power unit. Also, allows you to configure Ethernet-connected UPS devices, and specify how the NAS server or cluster reacts to power failures.

Procedure

1. Navigate to **Home > Server Settings > UPS Configuration** to display the **UPS Configuration** page.

At the bottom of the page is the UPS Devices table listing the configured UPS devices. For each UPS device, the table lists its IP address, percentage charged, runtime remaining on the UPS, and online status (whether currently supplying power).

Checking SMU status

Procedure

1. Navigate to **Home > SMU Administration > SMU Status** to display the **SMU Status** page.

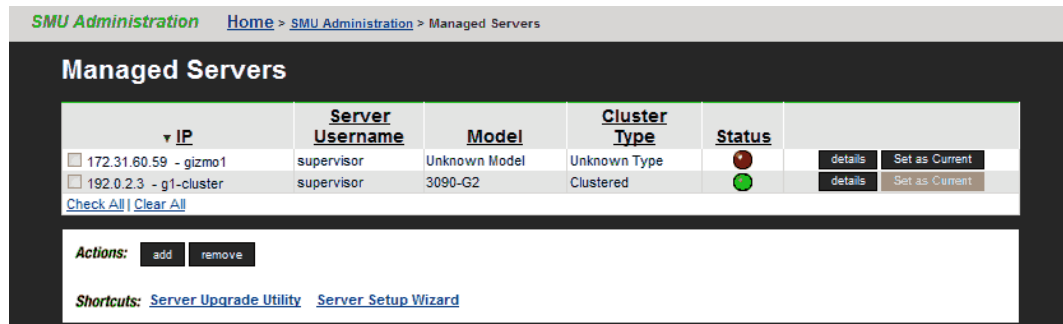
Field/Item	Description
Services	Quorum Service v2. For server firmware, version 10.0 or later. Used by a cluster, which has become partitioned by a network failure, to determine which partition is allowed to talk to the storage.

Field/Item	Description
	<p>Quorum Device. For server firmware, version 8.1 or earlier. Used by a cluster which has become partitioned by a network failure, to determine which partition is allowed to talk to the storage.</p> <p>Database. Allows communication between the SMU and the servers.</p>
Status	The desired state of these services is OK, and if a service is not running correctly, an error message is displayed.
Actions	<p>Displays the status of the SMU operating system. This is the actual output gathered from the Unix <code>top</code> command, and indicates the current running status of the SMU's internal processes.</p> <p>If the database service status is not running correctly an error message is displayed in the Status column. You can restart the database service by clicking restart.</p>
Top	Displays the status of the SMU operating system. This is the actual output gathered from the Unix <code>top</code> command, and indicates the current running status of the SMU's internal processes.
SMU Disk Usage (df)	Displays the details of the space used in each of the partitions of the SMU hard disk. This is the actual output gathered from the Unix <code>df</code> command.

Monitoring multiple servers

Procedure

1. Navigate to **Home > SMU Administrator > Managed Servers** to display the **Managed Servers** page.



Field/Item	Description
IP	IP address of the server. This should be the Administration Services IP address, as used on the private management network (for example, 192.0.2.x).
Server Username	User name of the NAS server.
Model	Displays the NAS server model number. For a cluster with different server models, this field displays "mixed", and the specific server models can be displayed in the Cluster Configuration page.
Cluster Type	Cluster type (for example, Node or Clustered).
Status	The color indicates the current status of the server: <ul style="list-style-type: none"> Green indicates that the server is operating normally (not showing an alert condition). Amber indicates a warning (operating normally, however, action should be taken to maintain normal operation). Red indicates a critical condition (the server is no longer functioning properly).
details	Opens Modify Managed Server page, which contains detailed information about contacting or managing the server.
Set as Current	Makes the currently selected server or cluster the currently managed server/cluster.
add	Adds a server or cluster that will then be managed by this SMU.
remove	Removes one or more selected servers or clusters. When a server or cluster is removed: <ul style="list-style-type: none"> Replication policies and schedules are deleted.

Field/Item	Description
	<ul style="list-style-type: none"> • Data migration policies and schedules are deleted. • The system monitor for that server is deleted. • Racks managed by that server are deleted.
Server Upgrade Utility	Opens the Server Upgrade Utility .
Server Setup Wizard	Opens the Server Setup Wizard .

When a server is removed:

- Replication policies and schedules are deleted.
- Data migration policies and schedules are deleted.
- The system monitor for that server is deleted.
- Racks managed by that server are deleted.

Monitoring storage subsystems with Hitachi Device Manager

The Hitachi Device Manager (HDvM) can be used to monitor HDS storage subsystems attached to Hitachi NAS Platform managed by an SMU. This functionality is enabled using Web Manager to configure a connection from the SMU to the HDvM server, then using the HDvM GUI to configure the HDvM to display information about the HDS storage subsystems attached to the servers managed by the SMU. (For information about configuring or using the Hitachi Device Manager, refer to your Hitachi Device Manager documentation, or contact Hitachi Data Systems.)

After an HDvM server has been specified, the SMU provides HDvM with information about the system drives (SDs) on the HDS storage subsystems attached to the NAS Platforms managed by the SMU.

Managing HDvM server connections

To define the HDS storage subsystems about which the SMU provides information to the HDvM, you specify which servers are connected to those storage subsystems. You can specify which HDvM servers are provided with information, and you can control which Hitachi NAS Platforms are included in the information provided to each HDvM server. This section covers specifying which HDvM servers the SMU provides with information, and defining the connection details for HDvM servers.



Note: The HDvM is provided with information about all SDs on any HDS storage subsystems attached to a Hitachi NAS Platform, if the SMU can collect this information from the Hitachi NAS Platforms.

To display HDvM servers that have been specified, and to see the Hitachi NAS Platforms about which the SMU is providing SD-related information to an HDvM server:

Procedure

1. Navigate to **Home > Storage Management > Hitachi Device Managers** to display the **Hitachi Device Managers** page.
 - If one or more HDvM servers have been specified, a list of defined HDvM servers is shown, displaying the host name or IP address of each HDvM server, and the date and time the last update was provided to that HDvM server. Click **details** to display the **Device Manager** details page that includes details about the HDvM server connection, and which servers are monitored by HDvM.

Storage Management Home > Storage Management > Hitachi Device Managers

Hitachi Device Managers

HDvM Host Name / IPAddress	Last Updated
<input type="checkbox"/> 192.168.42.121	

[Check All](#) | [Clear All](#)

Actions: [sync now](#) [remove](#) | [add](#)

- If an HDvM server has not yet been specified, the **Hitachi Device Managers** list is empty

Connecting the SMU to an HDvM server

Specifying connection details allows the SMU to communicate with the HDvM server and send information to HDvM.

Procedure

1. Navigate to **Home > Storage Management > Hitachi Device Managers** to display the **Hitachi Device Managers** page.
2. Click **add** to display the **Add Device Manager** page.

Storage Management Home > Storage Management > Hitachi Device Managers > Add Device Manager

Add Device Manager

Host Name / IP Address:

Username:

Password:

Use port for sending HDvM updates.

Monitored Servers

Available Servers

All Servers
gizmo1 (172.31.60.59)
g1-cluster (192.0.2.3)

Selected Servers

[OK](#) [cancel](#)

3. Specify HDvM server host name or IP address, valid HDvM user account name and password, and port for communicating with the HDvM server.

4. From the **Available Servers** list, select the servers with HDS storage subsystems to be monitored through HDvM.
5. Click **OK** to update the SMU's list of connections.



Note: The SMU sends information about HDS storage subsystems attached to the Hitachi NAS Platforms to the HDvM at 3:00 AM every morning.

6. Using the Hitachi Device Manager's GUI, add the storage subsystems attached to the Hitachi NAS Platforms managed by the SMU to the list of managed storage subsystems.
For information about configuring or using the Hitachi Device Manager, refer to your Hitachi Device Manager documentation, or contact Hitachi Data Systems.

Changing HDvM server connection details

Changing connection details allows the you to control the information the SMU sends to HDvM. You can specify which HDvM server the SMU connects to, and you can define which Hitachi NAS Platforms are included in the information sent to the HDvM.

Procedure

1. Navigate to **Home > Storage Management > Hitachi Device Managers** to display the **Hitachi Device Managers** page.
2. Click **details** to display the **Device Manager** details page.

The screenshot shows the 'Device Manager' configuration window. At the top, the breadcrumb navigation is 'Storage Management > Home > Storage Management > Hitachi Device Managers > Device Manager'. The window title is 'Device Manager'. It contains several input fields: 'Host Name / IP Address' with the value '192.168.42.121', 'Username' with 'hds', 'Password' with masked characters, and 'Use port' with '2001' and a note 'for sending HDvM updates.'. Below these is a section titled 'Monitored Servers' which is divided into two panes. The 'Available Servers' pane on the left contains 'All Servers' and 'g1-cluster (192.0.2.3)'. The 'Selected Servers' pane on the right contains 'gizmo1 (172.31.60.59)'. At the bottom of the panes are 'OK' and 'cancel' buttons.

3. Modify the HDvM server host name or IP address, valid HDvM user account name and password, and port for communicating with the HDvM server, as needed.
4. From the **Available Servers** list, select the servers with HDS storage subsystems to be monitored through HDvM, and add to the **Selected Servers** list.



Note: If you remove a managed server from the managed servers list after it has been monitored through HDvM, you must also manually remove it from the **Host** list in the HDvM GUI

5. Click **OK** to save the changes.
-



Note: Changes to the HDvM server connection details are effective immediately, but changes to which Hitachi NAS Platforms are to be monitored become effective the next time information is sent to the HDvM. (At 3:00 AM the next morning.)

Removing HDvM server connections

Procedure

1. Navigate to **Home > Storage Management > Hitachi Device Managers** to display the **Hitachi Device Managers** page.
2. Fill the check box next to the HDvM server connection you want to remove, and click **remove** to delete the connection.

Performance graphs


Data about system performance, load, and capacity is gathered at short intervals (typically every 10 seconds) and these data points are kept for the previous 24 hours. For data older than 24 hours, the data is periodically aggregated; the raw data is averaged and only a single hourly average value is kept. Aggregating the data in this fashion minimizes the overhead of millions of data points, and allows data to be retained long-term storage. If all the data points were retained, the data set would quickly grow, and would become unmanageable. The system stores the performance data for one year in order to provide a long-term, historical view of system performance.

- ☐ [Available performance graphs](#)
- ☐ [Controlling the performance graph display](#)
- ☐ [Storage server statistics](#)
- ☐ [Event logging and notification](#)
- ☐ [File system auditing](#)
- ☐ [FTP auditing](#)
- ☐ [Monitoring Fibre Channel switches](#)

Available performance graphs

The data can be displayed on graphs for review. The following graphs are available:

- File system performance (operations per second)
- File system capacity
- Storage pool capacity
- Node operations (a stand-alone server is displayed as a single node). The following graphs are available:

Graph	Describes	Units of measurement
File System Ops/Sec	For each of the selected file systems (up to five), the number of operations the selected file systems is processing, either for read or write.	Operations/second
File System Capacity	The total capacity of a single file system, including live data and snapshot usage.	MB (megabytes), GB (gigabytes), TB (terabytes), or PB (petabytes)  Note: The maximum size of a WFS-2 file system is 1 PB, but a 1 PB file system is <i>only</i> supported on an HDP storage pool.
Storage Pool Capacity	The total capacity of a single storage pool.	MB (megabytes), GB (gigabytes), or TB (terabytes)
Node ops	Operations per second on the node.	Operations/second
Ethernet Throughput	Ethernet throughput; both RX (received) and TX (transmitted).	Mb/sec (megabits per second)
System load	File system load on the hardware of the server/node. Refer to the hardware references for information about the hardware in your server/node.	% (percentage)
Disk latency	Disk read and write latency, and disk stripe write latency.	ms (milliseconds)
Fibre Channel throughput	Fibre Channel throughput; both RX (received) and TX (transmitted).	Mb/sec (megabits per second)

Graph	Describes	Units of measurement
Cache and heap usage	FSI cache and heap usage.	% (percentage)
NVRAM waited allocs	NVRAM waited allocations	Number (each)
Running Bossock Fibers	Running Bossock Fibers	Number (each)

Controlling the performance graph display

When displaying performance and capacity data, you can specify the date range to be displayed by the graph, either selecting one of the built-in ranges, or you can specify the date range for the display (a custom display range). To specify one of the standard display ranges, click one of the links in the display control above the upper right corner of the graph:

10m [1h](#) [1d](#) [1w](#) [1m](#) [3m](#) [1y](#) [Custom](#) [Download](#)

The standard display ranges include the following:

- 10m displays the last 10 minute's worth of collected data.
- 1h displays the last hour's worth of collected data.
- 1d displays the last day's (the last 24 hours, not the last calendar day) worth of collected data.
- 1w displays the last week's worth of collected data.
- 1m displays the last month's worth of collected data.
- 3m displays the last three months's worth of collected data.
- 1y displays the last year's worth of collected data.



Note: Because data older than 24 hours is periodically aggregated into hourly averages, date ranges that are longer than 24 hours will display aggregated data for periods more than 24 hours in the past.

You can specify a custom date range for the display of the performance graph by clicking Custom.

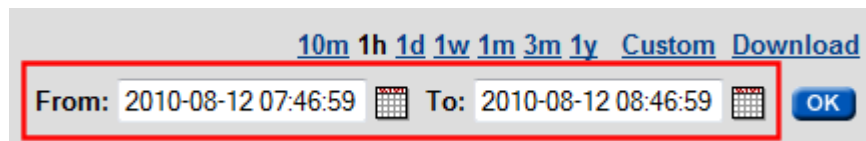
You can download data for later use by clicking Download.

Displaying a custom date range

Specifying a custom date range for the display of the performance graph allows you to examine system behavior under certain specific periods of time, which allows you to assess the impact of changes to the system or changes in the load on the system. To display a custom date range:

Procedure

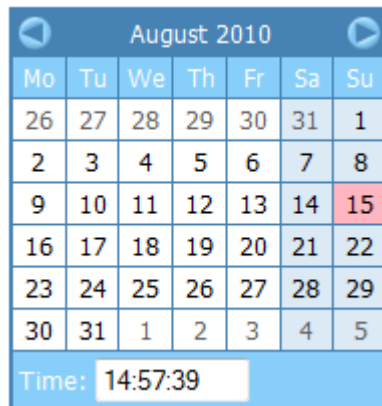
1. From the **Status and Monitoring** page, display the performance graph containing the data of interest.
 - Click **Performance Graphs** to display the **Performance Graphs** page. You can then display a custom date range for all the inset graphs on the page, or click an inset graph to access one of the following full-size graphs on its own page:
 - Node ops
 - Ethernet throughput
 - System load
 - Disk latency
 - Fibre Channel throughput
 - Cache and heap usage
 - NVRAM waited allocs
 - Running Bossock fibers
 - Click **File System Ops/Sec** to access the **File System Ops/Sec** page.
 - Click **File System Capacity** to access the **File System Capacity** page.
 - Click **Storage Pool Capacity** to access the **Storage Pool Capacity** page.
2. After displaying the performance graph containing the data of interest, click **Custom** to display the custom date range controls.



10m 1h 1d 1w 1m 3m 1y Custom Download

From: 2010-08-12 07:46:59 To: 2010-08-12 08:46:59 OK

3. Click the **From** calendar icon to display the calendar control to select the start of the date range.



The following table describes the calendar control:

Item	Description
Left arrow	Click the left arrow to change to the month preceding the month displayed in the control.
Right arrow	Click the right arrow to change to the month following the month displayed in the control.
Day selector	Click a day in the calendar control to select that day. Note that selecting a day returns you to the graph page.
Time	In the Time edit box, you can specify the time, using the 24-hour notation and in the format <i>hh:mm:ss</i> .

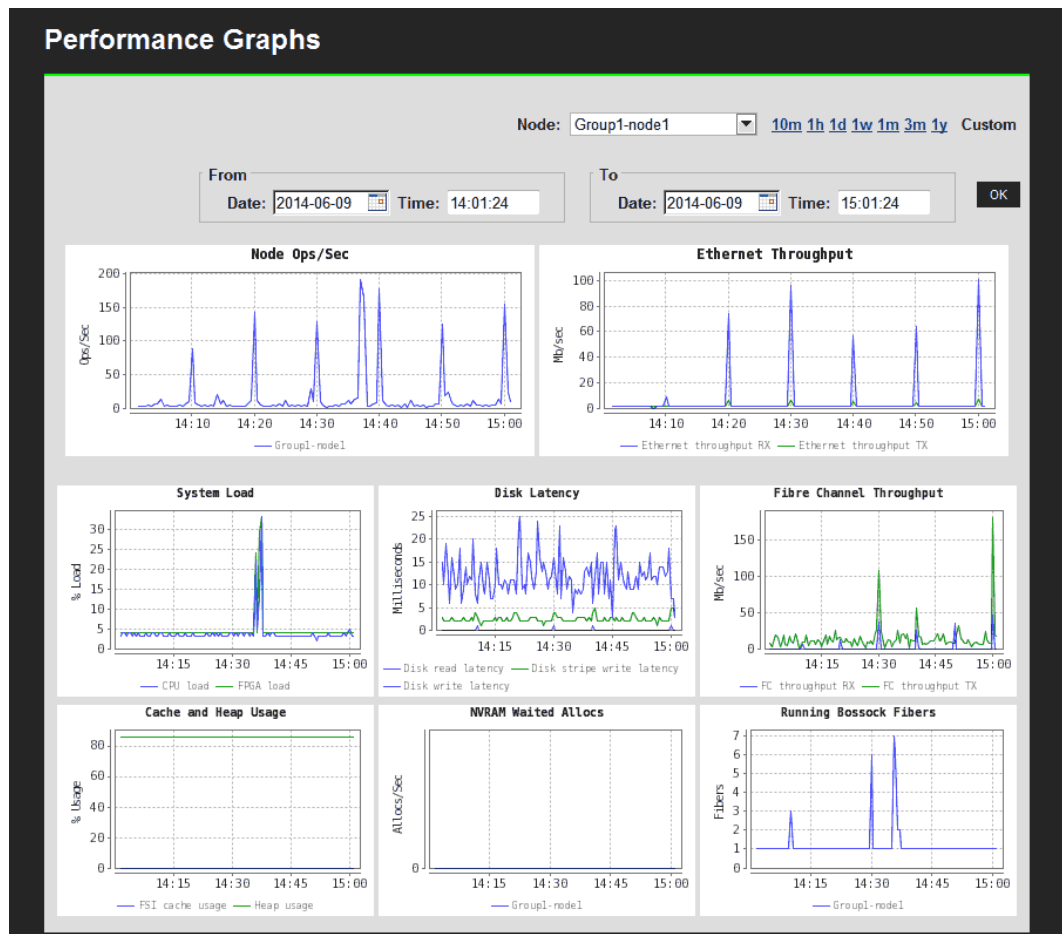
4. Click the **To** calendar icon to display the calendar control, which you then use to select the end of the date range.
5. Click **OK** to display the performance graph using data from the specified date range.
6. Optionally, click **Download** to download the data for later use.

Displaying the Performance Graphs page

The **Performance Graph** page provides an overview of a single node's performance status.

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs** to display the **Performance Graphs** page.



Item	Description
Node	<p>List includes all nodes in the cluster. For a stand-alone server, only one node will appear in the list.</p> <p>By default, the graphs on this page display system performance data for the first cluster node of the cluster; however, you can select a specific node from the drop down list in order to view its unique performance data.</p> <p>At any time, to display the performance information for a node other than the one currently being displayed, select the node from the list.</p>
Date Range Display Controls	Allows you to select the date range for the graphs currently being displayed.
Download (link)	Downloads data for later use by clicking Download .

Item	Description
Performance Graphs	<p>The individual performance graphs on this page are links to pages that display a full-size version of that same graph. To display the page containing the full-sized graph, click the graph you want to display. The following graphs are displayed on this page:</p> <ul style="list-style-type: none"> • Node Ops/Sec • Ethernet Throughput • System Load • Disk Latency • Fibre Channel Throughput • Cache and Heap Usage • NVRAM Waited Allocs • Running Bossock Fibers

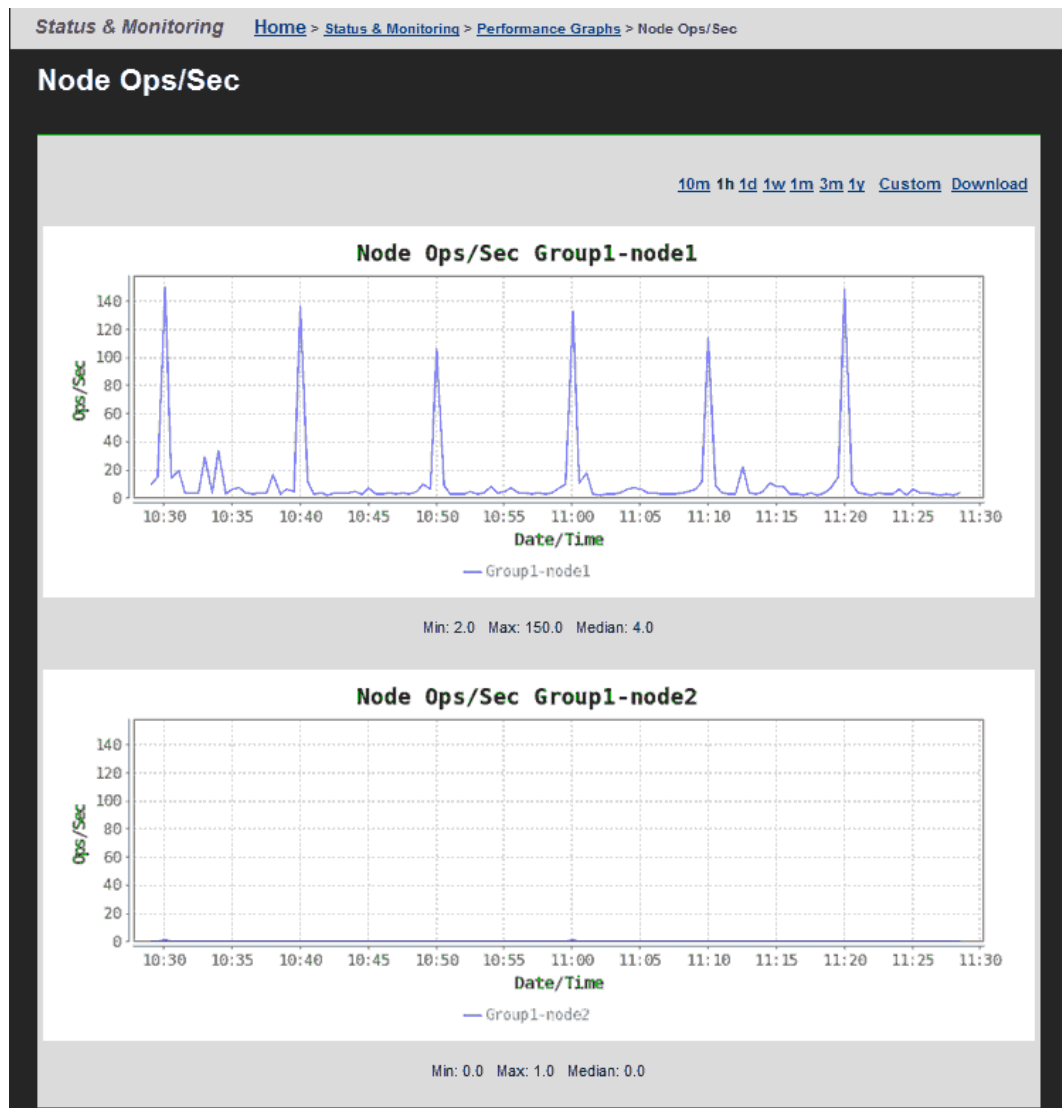
2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying Node Ops/Sec

The **Node Ops/Sec** page displays system performance data for node ops/sec for all the cluster nodes of the currently managed cluster. For example, if the server has four cluster nodes, four graphs for node ops/sec are displayed on this page.

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > Node Ops/Sec** to display the **Node Ops/Sec** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.

Item	Description
Ops/Sec	The number of recorded operations per second.
Date/Time	The currently selected date/time range.
Min/Max/Median	The minimum, maximum, and median number of operations per second detected in the currently displayed date range.

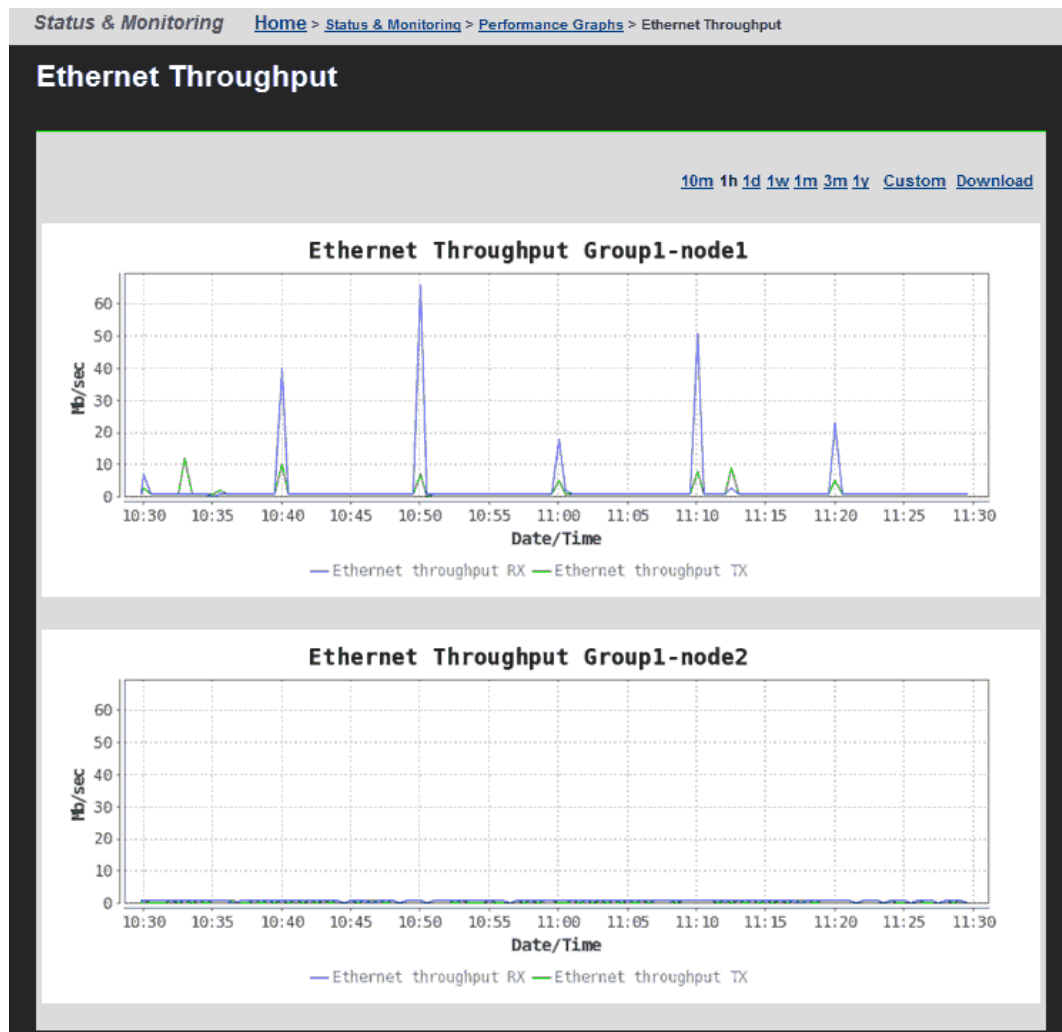
2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying Ethernet Throughput

The **Ethernet Throughput** page displays system performance data for Ethernet throughput, for both transmission and reception, a per-node basis for all the cluster nodes of the currently managed cluster.

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > Ethernet Throughput** to display the **Ethernet Throughput** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
Mb/Gb/sec	The number of megabits (Mb) or gigabits (Gb) transmitted (TX) and received (RX) per second.
Date/Time	The currently selected date/time range.

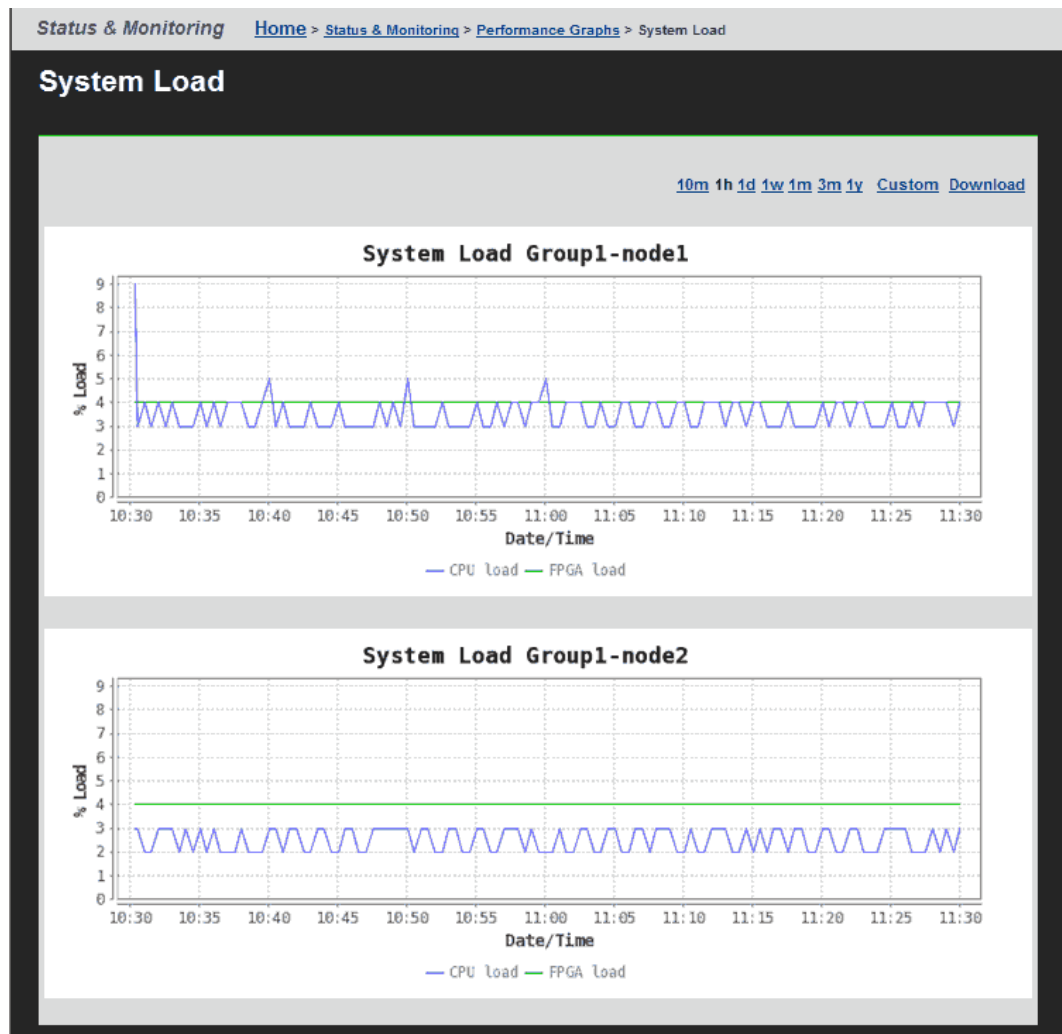
2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying System Load

The **System Load** page displays system performance data for system load for all the cluster nodes of the currently managed cluster. The graph displays the percentage of maximum usage at a given point in time.

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > System Load** to display the **System Load** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
% Load	For servers or nodes containing FSB and FSA module, file system loads are displayed as a percentage of maximum projected capacity of those modules.

Item	Description
	For servers and nodes containing MFB and MMB boards, file system loads are displayed as a percentage of the maximum projected capacity of those boards. Refer to the hardware reference for your system for more information on the hardware in your server/nodes.
Date/Time	The currently selected date/time range.

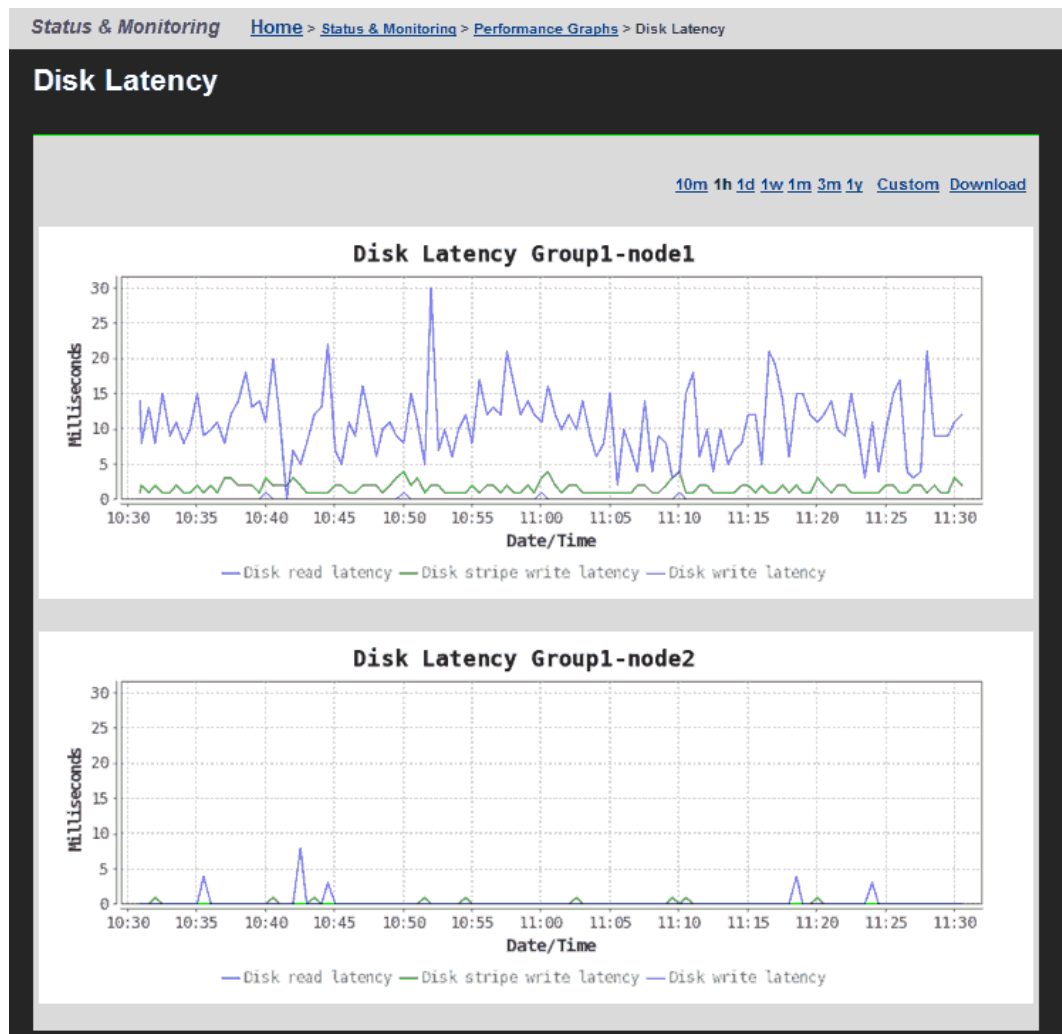
2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying Disk Latency

The **Disk Latency** page displays system performance data for disk latency (disk read, write, and disk stripe write) in milliseconds for all the cluster nodes of the currently managed cluster.

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > Disk Latency** to display the **Disk Latency** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
Milliseconds	The number of milliseconds of latency for disk read operations, write operations, and disk stripe write operations.
Date/Time	The currently selected date/time range.

2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying Fibre Channel Throughput

The **Fibre Channel Throughput** page displays system performance data for Fibre Channel throughput (transmission and reception) for all the cluster nodes of the currently managed cluster in megabits per second (Mb/sec).

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > Fibre Channel Throughput** to display the **Fibre Channel Throughput** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.

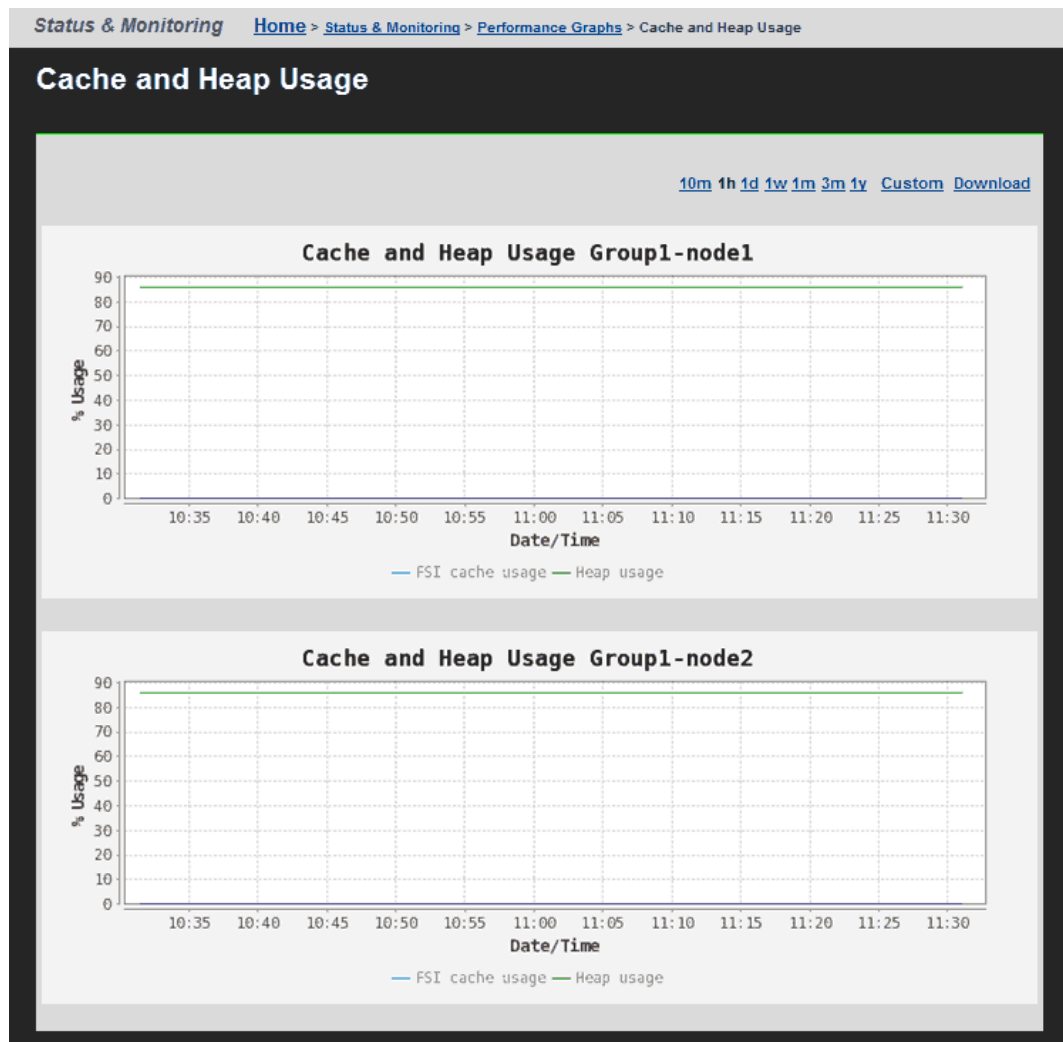
Item	Description
Mb/Gb/sec	The aggregated number of megabits (Mb) or gigabits (Gb) transmitted (TX) and received (RX) per second through the Fibre Channel ports of the node.
Date/Time	The currently selected date/time range.

2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying Cache and Heap Usage

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > Cache and Heap Usage** to display the **Cache and Heap Usage** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
% Usage	The percentage of the total available FSI cache and heap in use.
Date/Time	The currently selected date/time range.

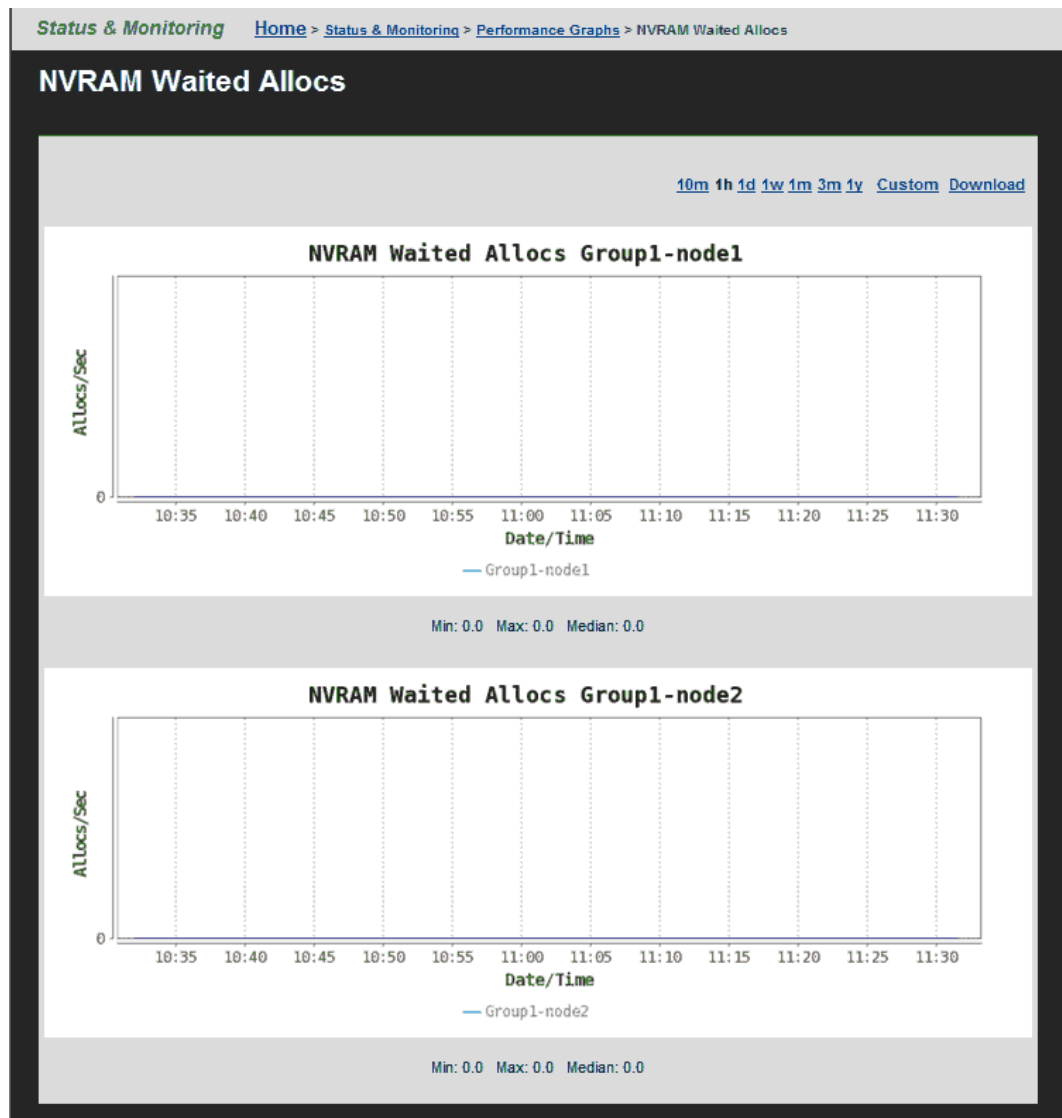
2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying NVRAM Waited Allocs

The **NVRAM Waited Allocs** page displays system performance data for the number of NVRAM waited allocs per second for all the cluster nodes of the currently managed cluster. An NVRAM waited alloc indicates that a file system has had to wait for NVRAM space to be allocated, which means that performance has been negatively impacted.

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > NVRAM Waited Allocs** to display the **NVRAM Waited Allocs** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.

Item	Description
Allocs/Sec	The average number of times per second a node's NVRAM had to wait for NVRAM space to be allocated.
Date/Time	The currently selected date/time range.
Min/Max/Median	The minimum, maximum, and median number of NVRAM waited allocs per second detected in the currently displayed date range.

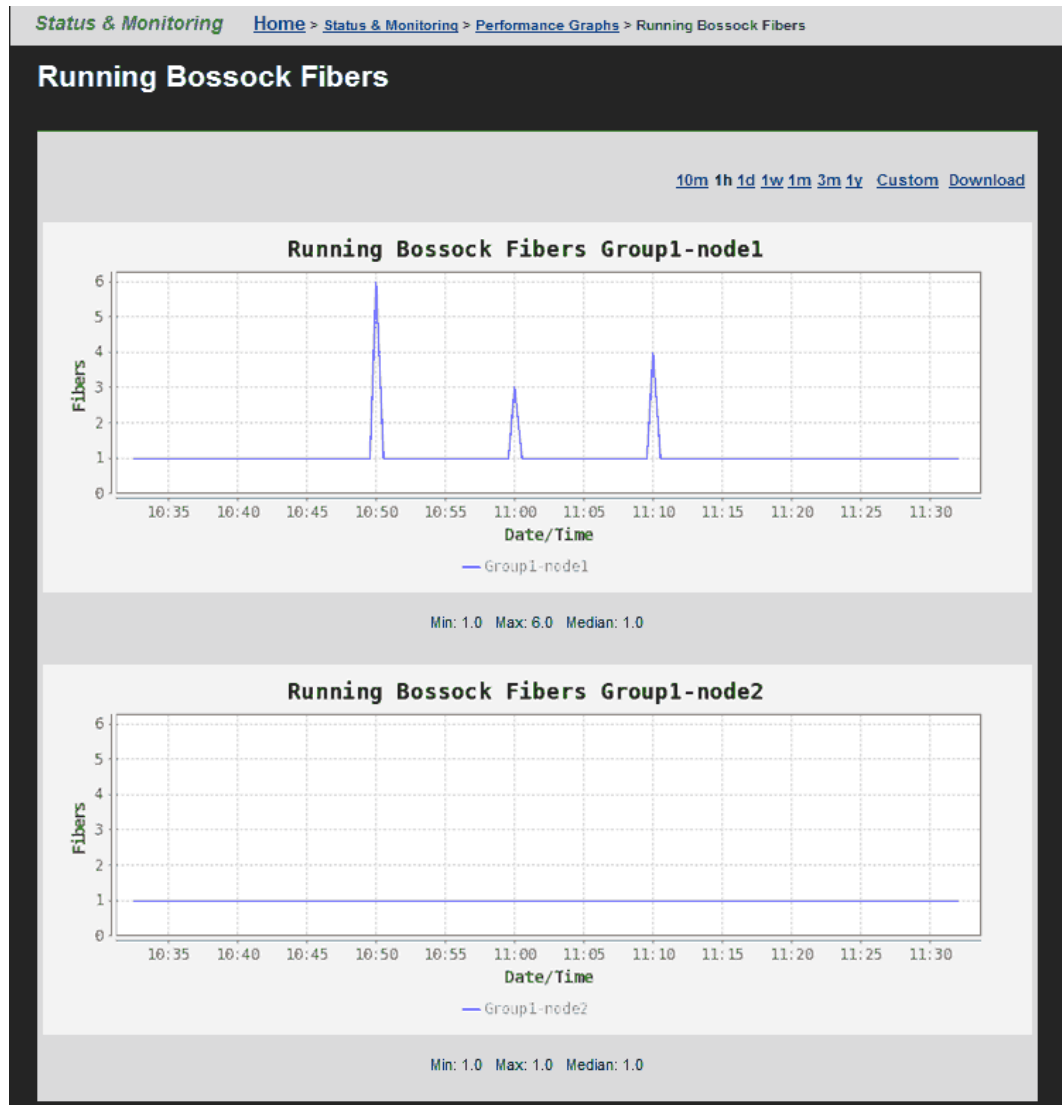
2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying Running Bossock Fibers

The **Running Bossock Fibers** page displays system performance data for the number of bossock fibers in use for all the cluster nodes of the currently managed cluster. A bossock fiber is a network receive thread, and most file serving traffic received from the network is handled synchronously by one of a pool of these threads.

Procedure

1. Navigate to **Home > Status & Monitoring > Performance Graphs > Running Bossock Fibers** to display the **Running Bossock Fibers** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.

Item	Description
Fibers	The number of bossock fibers currently in use.
Date/Time	The currently selected date/time range.
Min/Max/Median	The minimum, maximum, and median number of bossock fibers in use during the currently displayed date range.

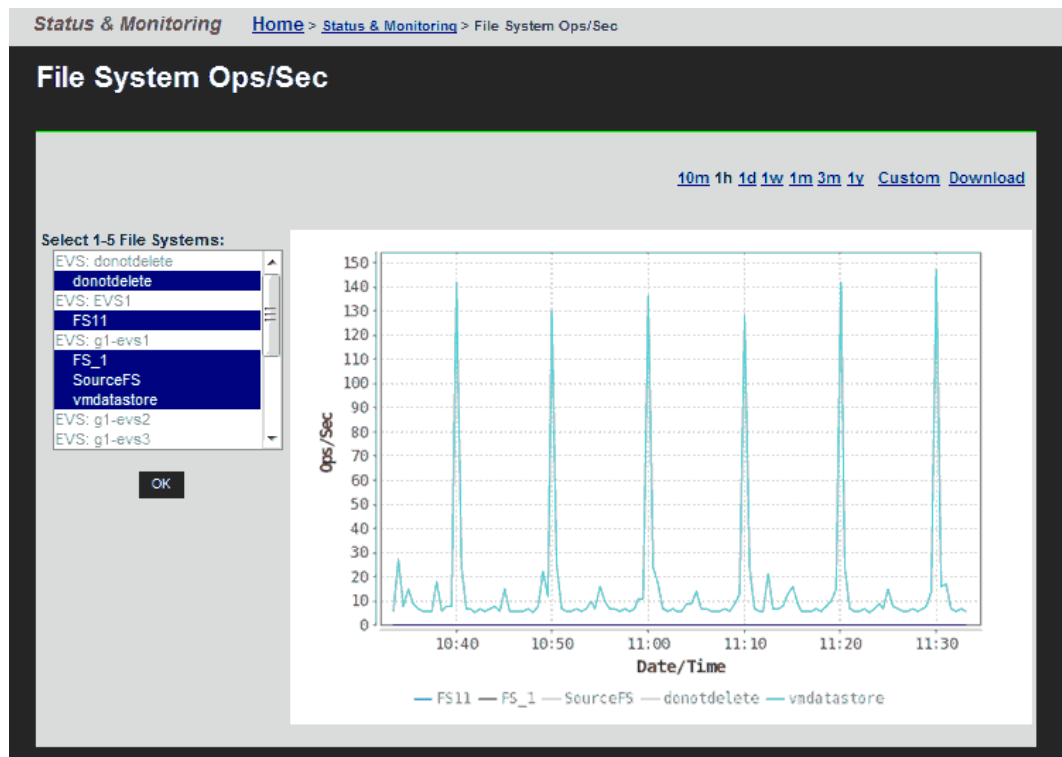
2. Optionally, change the date range displayed in the graph using the date range controls.

Displaying File System Ops/Sec

The **File System Ops/Sec** page displays file system performance data for up to five selected file systems of the currently managed cluster or server. This graph displays the average number of read and/or write operations performed per second by the selected file systems.

Procedure

1. Navigate to **Home > Status & Monitoring > File System Ops/Sec** to display the **File System Ops/Sec** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
Select 1-5 File Systems	Displays all the file systems in the currently managed server or cluster. To select file systems, and display the performance data for that file system: <ol style="list-style-type: none"> 1. On your keyboard, press and hold the Control (Ctrl) key. 2. Highlight up to five file systems. 3. Click OK.
Ops/Sec	The number of recorded operations per second.
Date/Time	The currently selected date/time range.

2. Using the **Select a File System** list, select the file systems, and display the performance data.

3. Optionally, change the date range displayed in the graph using the date range controls.

Displaying File System Capacity


The **File System Capacity** page displays file system capacity data for up to five selected file systems of the currently managed cluster or server. This graph displays the total space used, in MB (megabytes), GB (gigabytes) or TB (terabytes), including live data and snapshot usage by the selected file systems.



Note: The maximum size of a WFS-2 file system is 1 PB, but a 1 PB file system is *only* supported on an HDP storage pool.

Procedure

1. Navigate to **Home > Status & Monitoring > File System Capacity** to display the **File System Ops/Sec** page.



Item	Description
1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
Select a File System	Displays all the file systems in the currently managed server or cluster. To select file systems, and display the performance data for that file system: <ol style="list-style-type: none"> 1. On your keyboard, press and hold the Control (Ctrl) key. 2. Highlight up to five file systems. 3. Click OK.
MB, GB, or TB	The total space used, in MB (megabytes), GB (gigabytes) or TB (terabytes), including live data and snapshot usage by the selected file systems.

Item	Description
Date/Time	The currently selected date/time range.


2. Using the **Select a File System** list, select the file systems, and display the performance data for the selected file systems.
3. Optionally, change the date range displayed in the graph using the date range controls.

Displaying Storage Pool Capacity

The **Storage Pool Capacity** page displays capacity and usage data for the selected storage pool (span) of the currently managed cluster or server. This graph displays the total space allocated and used, in MB (megabytes), GB (gigabytes) or TB (terabytes), by all file systems in the selected storage pool.

Procedure

1. Navigate to **Home > Status & Monitoring > Storage Pool Capacity** to display the **Storage Pool Capacity** page.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
Select a Storage Pool	Displays all the storage pools (spans) in the currently managed server or cluster. Highlight the storage pool for which you want to display data, and click OK .
MB, GB, or TB	The storage pool capacity (the total space allocated to the storage pool) and the amount of storage in use by the storage pool. The capacity and usage are indicated in MB (megabytes), GB (gigabytes), or TB (terabytes).
Date/Time	The currently selected date/time range.

2. Using the **Select a Storage Pool** list, select the storage pool, and display the performance data for that storage pool.

3. Optionally, change the date range displayed in the graph using the date range controls.

Downloading performance data

You can download performance and capacity data from any of the performance or capacity graphs. The data is downloaded to your client in a compressed .CSV (comma-separated values) text file, that can be decompressed and then opened by many applications. The data downloaded is from the currently specified date range, and it includes the following data points:

- For all data collected within the last 24 hours, all data points collected at the minimum collection interval for the information shown in the graph. The minimum collection interval differs, and depends on the data being collected.
- For all data collected more than 24 hours in the past, data points are aggregated and averaged into hourly values, and only the hourly values are retained for up to one year.

Procedure

1. Navigate to the graph containing the data you want to download:
 - Click **Performance Graphs** to display the **Performance Graphs** page. You can then display a custom date range for all the inset graphs on the page, or click an inset graph to access one of the following full-size graphs on its own page:
 - Node ops
 - Ethernet throughput
 - System load
 - Disk latency
 - Fibre Channel throughput
 - Cache and heap usage
 - NVRAM waited allocs
 - Running Bossock fibers
 - Click **File System Ops/Sec** to access the **File System Ops/Sec** page.
 - Click **File System Capacity** to access the **File System Capacity** page.
 - Click **Storage Pool Capacity** to access the **Storage Pool Capacity** page.
2. Click **Custom** to set the date range for the data you want to download
3. Click **Download** (located next to the date range display controls) to display a dialog that allows you to open the file or to save the file in a destination director. The dialog that appears depends on your browser, but in general you should indicate that you want to save the file, and click **OK** or **Save**.

4. Specify (or navigate to) the destination directory, and then click **OK** or **Save** to save the downloaded file in that directory.
The download process might take anywhere from a few seconds to a few minutes, depending on how much data is in the file to be downloaded. After the file has been downloaded, it can be decompressed using any utility that can decompress zip files.

Storage server statistics

The Hitachi NAS Platform provides extensive statistics that can be used to monitor operation. These statistics include:

- Networking (Ethernet and TCP/IP)
- Fibre Channel
- File access protocols (CIFS, NFS, and FTP)
- Block access protocols (iSCSI)
- Management access for supported access protocols (SNMP and SSC for some NAS server models, or SNMP, SSC, Telnet, and SNMP for more recent NAS server models)
- Virus scanning

Data is gathered at short intervals (typically every 10 seconds) and these data points are kept for the previous 24 hours. For data older than 24 hours, the data is periodically aggregated; the raw data is averaged and only a single hourly average value is kept. Aggregating the data in this fashion minimizes the overhead of millions of data points, and allows data to be retained long-term storage. If all the data points were retained, the data set would quickly grow, and would become unmanageable. The system stores the performance data for 1 year in order to provide a long-term, historical view of system performance.

In addition to the statistics, performance graphs are also provided.

Network statistics

Fibre Channel, Ethernet, and TCP/IP statistics for the server (per port in 10-second timeslices) are available. These statistics pages display activity since the previous reboot, or since the point when statistics were last reset.

Ethernet statistics

Ethernet statistics (per port in 10-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Displaying Ethernet Statistics

Procedure

1. Navigate to **Home > Status & Monitoring > Ethernet Statistics** to display the **Ethernet Statistics** page.

Status & Monitoring Home > Status & Monitoring > Ethernet Statistics		
Ethernet Statistics		
Cluster Node: Group1-node1		
change...		
Last Reset: 2014-05-28 17:38:13 (UTC-0700) reset Last Refreshed: 2014-06-13 11:35:14 (UTC-0700)		
	Receive Rate (bytes/second)	Transmit Rate (bytes/second)
Instantaneous	8,791	1,090
Peak	49,857,438	51,328,926
	Transmitted OK	Received OK
Bytes	73,606,700,081	309,644,161,128
Packets	176,733,574	239,637,319
	Total	
Bytes	383,250,861,209	
Packets	416,370,893	
	Receive Errors	Transmit Errors
Packet drops	0	0
CRC errors	0	-
Oversized packets	0	-
Fragmented packets	0	-
Collisions	0	-
Jabbers	0	-
Undersized packets	0	-
Unknown Protocol	1831298	-
One collision	-	0
Multiple collisions	-	0
Excessive collisions	-	0
Late collisions	-	0

Field/Item	Description
Cluster Node	When connected to a NAS server cluster, indicates the node for which the statistics are displayed.
change	Opens the Select a Cluster Node page in which you can select a different node for which to display statistics.
Last Reset	Displays the date and time the statistics on this page were reset.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Receive Rate	The amount of data received in bytes per second. Includes the current (Instantaneous) and Peak throughput.
Transmit Rate	The amount of data transmitted in bytes per second. Includes the current (Instantaneous) and Peak throughput.
Transmitted OK	The total number of Bytes and Packets successfully Transmitted.
Received OK	The total number of Bytes and Packets successfully Received.
Total	The total number of Bytes and Packets Transmitted and Received.
Errors	Lists the number of Receive Errors and Transmit Errors logged on the server/node. The following types of errors are reported. For the five

Field/Item	Description
	<p>collision errors, the values will always be 0 for the file-serving ports (ge, tg, and ag), as only full duplex point-point connections are supported.</p> <ul style="list-style-type: none"> • Packet drops • CRC errors • Oversized packets • Collisions • Jabbers • Undersized packets • Unknown protocol • One collision • Multiple collisions • Excessive collisions • Late collisions
reset	Enables you to reset the statistics when needed.

Displaying aggregated ports or per-port Ethernet statistics

Procedure

1. Navigate to **Home > Status & Monitoring > Ethernet Statistics (per port)** to display the **Aggregated Ports Ethernet Statistics** page.
2. Click **Physical Ports Ethernet Statistics** to display the statistics for individual ports.

The following table describes the fields in both of these pages:

Field/Item	Description
Cluster Nodes	When connected to a NAS server cluster, this field indicates the node for which the statistics are displayed. To display statistics for another node, click the change button.
change	Displays the Select a Cluster Node page in which you can select a different node for which to display statistics.
Last Refreshed	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Bytes	Displays the number of bytes Transmitted OK and Received OK, and the Total number of bytes.
Packets	<p>Displays the number of packets Transmitted OK and Received OK, and the Total number of packets. This also displays:</p> <ul style="list-style-type: none"> • Unicast Received • Broadcast Received • Multicast Received • Unicast Transmitted • Broadcast Transmitted • Multicast Transmitted

Field/Item	Description
Receive Throughput Rate	The receive rate in bytes/second for the Instantaneous (current) and Peak throughput.
Transmit Throughput Rate	The transmit rate in bytes/second for the Instantaneous (current) and Peak throughput.
Receive Errors	Lists the number of Receive Errors logged on the server/node. The following types of errors are reported: <ul style="list-style-type: none"> • Packet drops • CRC errors • Oversized packets • Fragmented packets • Collisions • Jabbers • Undersized packets • Unknown protocol
Transmit Errors	Lists the number of Transmit Errors logged on the server/node. The following types of errors are reported: <ul style="list-style-type: none"> • Packet drops • One collision • Multiple collisions • Excessive collisions • Late collisions
Link Status	Indicates the condition of the link. The possible values are either Up or Down.
MAC Addresses	Displays the MAC address of each port. For the aggregated statistics, the MAC address is for the port to which the aggregation is linked.
Last Reset Time	Displays the date and time when the statistics for this port were last reset to zero.
Select to Reset Statistics	Fill this check box for each port whose statistics you want to reset to zero. The statistics are reset when you click reset .
reset	Resets all statistics of the selected ports to zero.
Ethernet Statistics (per port) - Physical Ports	Goes to Ethernet Statistics (per port) Physical Ports page, where you can view the statistics for each of the defined ports in the HNAS server/cluster mode.

TCP/IP Statistics

The TCP/IP statistics display activity since the last server reboot, or since the TCP/IP statistics were last reset. Both per-port and overall statistics are available on this page. The statistics are updated every 10 seconds.

Displaying TCP/IP statistics

Procedure

1. Navigate to **Home > Status & Monitoring > TCP/IP Statistics** to display the **TCP/IP Statistics** page.

Status & Monitoring [Home > Status & Monitoring > TCP/IP Statistics](#)

TCP/IP Statistics

Cluster Node: Group1-node1
[change...](#)

Last Reset: 2014-05-28 17:38:20 (UTC-0700) [reset](#) Last Refreshed: 2014-06-13 11:36:37 (UTC-0700)

TCP

TCP Connections

Currently Open: 13
Maximum Open: 25
Total Opened: 714390
Failed Connections: 66913

Packets

	TCP Packets	UDP Packets	ICMP Packets	ICMPv6 Packets	Other Packets
Transmitted	171190496	5183986	8	10	0
Received	234206051	1302263	172074	320	184511
Retransmitted	6252				
Invalid	0	0			184511
Unknown Port		385034			
Unknown Protocol					0

Field/Item	Description
Cluster Node	When the server is part of a cluster, this field identifies the node. Click change to change nodes, and view statistics for that node.
Last Reset	The date and time the statistics on this page were reset. Click reset to reset the statistics to zero.
Last Refreshed	The date and time this page was refreshed. The page automatically refreshes every 10 seconds.
TCP Connections	Displays statistics about the TCP connections. <ul style="list-style-type: none">• Currently Open is the number of currently open connections.• Maximum Open is the maximum number of connections opened at one time since the last reset.• Total Opened is the number of connections that have been opened since the last reset.• Failed Connections is the number of failed incoming and outgoing connections.
Packets	Lists the number of transmitted, received, retransmitted, invalid, unknown ports, and unknown protocols since the last reset for: <ul style="list-style-type: none">• TCP Packets• UDP Packets• ICMP Packets

Field/Item	Description
	<ul style="list-style-type: none"> • ICMPv6 Packets • Other Packets <p>An IP packet is invalid when any of the following is invalid:</p> <ul style="list-style-type: none"> • Header checksum • Length field (too long for the packet) • Source address • Destination address (this is the most common cause)

Displaying aggregated ports or per-port TCP/IP statistics

Procedure

1. Navigate to **Home > Status & Monitoring > TCP/IP Statistics (per port)** to display TCP/IP statistics for individual ports or aggregated for all ports.
2. Click **Physical Ports TCP/IP Statistics** to display the statistics for individual ports.

The following table describes the fields in both of these pages:

Field/Item	Description
Cluster Node	When connected to a NAS server cluster, indicates the node for which the statistics are displayed.
change	Opens the Select a Cluster Node page in which you can select a different node for which to display statistics.
Last Refreshed	Displays the date and time this page was refreshed. This page automatically refreshes every 10 seconds.
Bytes	Displays the number of bytes Transmitted OK and Received OK, and the Total number of bytes.
Packets	Displays the number of packets Transmitted OK and Received OK, and the Total number of packets.
Receive Throughput Rate	The receive rate in bytes/second for the Instantaneous (current) and Peak throughput.
Transmit Throughput Rate	The transmit rate in bytes/second for the Instantaneous (current) and Peak throughput.
Receive Errors	Lists the number of Receive Errors logged on the server/node. The following types of errors are reported: <ul style="list-style-type: none"> • Packet drops • CRC errors • Oversized packets • Fragments packets • Collisions • Jabbers • Undersized packets

Field/Item	Description
	<ul style="list-style-type: none"> Unknown Protocol
Transmit Errors	<p>Lists the number of Transmit Errors logged on the server/node. The following types of errors are reported:</p> <ul style="list-style-type: none"> Packet drops One collision Multiple collisions Excessive collisions Late collisions
Link Status	Indicates the condition of the link. The possible values are either Up or Down.
MAC Addresses	Displays the MAC address of each port. For the aggregated statistics, the MAC address is for the port to which the aggregation is linked.
Last Reset Time	Displays the date and time when the statistics for this port were last reset to zero.
Select to Reset Statistics	Selects which ports to reset.
reset	Resets the statistics.
TCP/IP Statistics (per port) - Aggregated Ports	Click to advance to the TCP/IP Statistics (per port) Aggregated Ports page.

Displaying TCP/IP detailed statistics

Procedure

1. Navigate to **Home > Status & Monitoring > TCP/IP Detailed Statistics** to display the **Detailed TCP/IP Statistics** page.

Status & Monitoring [Home](#) > [Status & Monitoring](#) > TCP/IP Detailed Statistics

TCP/IP Detailed Statistics

Cluster Node: Group1-node2
[change...](#)

Last Reset: 2014-07-22 13:41:17 (UTC-0700) [reset](#) Last Refreshed: 2014-07-25 10:17:55 (UTC-0700)

IP Errors	TCP Errors	UDP Errors
Invalid Header Field: 0	Invalid Checksum: 0	Short Packet: 0
Oversized Segment: 0		Invalid Checksum: 0
Invalid Source Address: 9402		
Invalid Option: 0		

Field/Item	Description
Cluster Node	When connected to a NAS server cluster, indicates the node for which the statistics are displayed.
change	Opens the Select a Cluster Node page in which you can select a different node for which to display statistics.
Last Reset	Displays the date and time the statistics on this page were reset.
reset	Resets the statistics to zero.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
IP Errors: Invalid Header Field	Displays the number of IP errors arising from an invalid header field.
IP Errors: Oversized Segment	Displays the number of fragmented TCP packets greater than the Maximum Transmission Unit (MTU) size when reassembled. The transmitting source made an error or the packet was corrupted in transit.
IP Errors: Invalid Source Address	Displays the number of IP packets with an invalid source address (often caused by DHCP broadcast requests using the source address 0).
IP Errors: Invalid Option	Displays the number of IP packets that were not decoded because the IP option length was invalid. The transmitting source made an error or the packet was corrupted in transit.
TCP Errors: Invalid Checksum	Displays the number of invalid TCP packet checksums. The transmitting source made an error or the packet was corrupted in transit.

Field/Item	Description
UDP Errors: Short Packet	Displays the number of UDP packets that were too short for the UDP header or length. The transmitting source made an error or the packet was corrupted in transit.
UDP Errors: Invalid Checksum	Displays the number of invalid UDP packet checksums. The transmitting source made an error or the packet was corrupted in transit.

Fibre Channel statistics

The Fibre Channel (FC) statistics for the server (per port in 10-second time slices) are available for activity, as of the previous reboot or the point when statistics were last reset.

Displaying Fibre Channel statistics

Procedure

1. Navigate to **Home > Storage Management > Fibre Channel Statistics** to display the **Fibre Channel Statistics** page.

Storage Management [Home](#) > [Storage Management](#) > [Fibre Channel Statistics](#)

Fibre Channel Statistics

Cluster Node: Group1-node1
[change...](#)

Last Reset: 2014-05-28 17:38:13 (UTC-0700) [reset](#) Last Refreshed: 2014-06-13 11:39:14 (UTC-0700)

Throughput		
	Receive Rate (bytes/second)	Transmit Rate (bytes/second)
Instantaneous	0	1,049,088
Peak	494,960,640	148,988,928
Total	380,503,902,208	3,087,730,749,440

I/O Requests				
	Disk Reads	Disk Writes	Tape Reads	Tape Writes
Total Requests	14,804,235	110,590,059	0	0
Total Responses	14,804,233	110,590,059	0	0

Total Requests & Responses
Requests: 453,859,320
Responses: 453,859,320

Cache
Hits: 414,695,258
Misses: 14,785,696

I/O Status Counters
Failed: 0
Resubmitted: 2

Total Errors
Loss of Signal: 0
Bad Receive Character: 58
Loss of Sync: 2
Link Fail: 2
Receive EOFa: 0
Discarded Frames: 0
Bad CRCs: 0
Protocol Errors: 0

Congestion
Instantaneous: 0
Peak over 24 hours: 0
Ave. over 24 hours: 0

Field/Item	Description
Cluster Nodes	When connected to a cluster, this field indicates the node for which the statistics are displayed. To display statistics for another node, click the change button.
change	Displays the Select a Cluster Node page, on which you can select a different node for which to display statistics.
Last Reset	Displays the date and time the statistics on this page were reset. To reset the statistics to zero, click the reset button.
Last Refreshed	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.

Field/Item	Description
Throughput	
Receive Rate	The amount of data received in bytes per second. Includes the Instantaneous (current), Peak, and Total throughput.
Transmit Rate	The amount of data transmitted in bytes per second. Includes the Instantaneous (current), Peak, and Total throughput.
I/O Requests	
Disk Reads	The number of read requests that the attached disk devices have received, and the number of responses sent.
Disk Writes	The number of write requests that the attached disk devices have received, and the number of responses sent.
Tape Reads	The number of read requests that the attached tape devices have received, and the number of responses sent.
Tape Writes	The number of write requests that the attached tape devices have received, and the number of responses sent.
Total Requests and Responses	The number of data requests that the server has received, and the number of responses it has sent out. These include both requests that have been sent to the storage devices and requests that the cache has served internally.
Cache	Number of hits (requests that the cache has served) and misses (requests not served by the cache and passed to the storage subsystem).
I/O Status Counters	Numbers of failed and resubmitted input and output requests.
Total Errors	Number of errors logged at the Fibre Channel interface. The following types of errors are reported: Loss of Signal, Bad Receive Character, Loss of Sync, Link Fail, Receive EOFa, Discarded Frames, BAD CRCs, and Protocol Errors.
Congestion	Displays congestion rates. Includes the Instantaneous (current), Peak (over the past 24 hours), and Average (average over the past 24 hours) rates.

Displaying per port Fibre Channel statistics

Procedure

1. Navigate to **Home > Storage Management > Fibre Channel Statistics (per port)** to display statistics for each of the defined ports.

Storage Management Home > Storage Management > Fibre Channel Statistics (per port)				
Fibre Channel Statistics (per port)				
Cluster Node: Group1-node1				
change...				
Last Reset: 2014-05-28 17:38:13 (UTC-0700) reset Last Refreshed: 2014-06-09 15:16:52 (UTC-0700)				
	FC 1	FC 2	FC 3	FC 4
Receive Throughput Rate (bytes/second)				
Instantaneous	0	0	0	0
Peak	249,005,056	0	253,886,464	0
Total	150,093,040,128	0	150,119,828,992	0
Transmit Throughput Rate (bytes/second)				
Instantaneous	896,000	0	313,856	0
Peak	75,357,696	0	81,975,296	0
Total	1,147,400,929,280	0	1,150,711,124,480	0
Total Errors				
Loss of Signal	0	0	0	0
Bad Receive Character	24	35	34	777
Loss of Sync	1	0	1	0
Link Fail	1	0	1	0
Receive EOFa	0	0	0	0
Discarded Frames	0	0	0	0
Bad CRCs	0	0	0	0
Protocol Errors	0	0	0	0
Congestion				
Instantaneous	0	0	0	0
Peak over 24 hours	0	0	0	0
Ave. over 24 hours	0	0	0	0
Speed	4 Gbps	4 Gbps	4 Gbps	4 Gbps
Type	N	N	N	N
Status	Up	Down	Up	Down
Enabled	Enabled	Disabled	Enabled	Disabled

Field/Item	Description
Cluster Nodes	When connected to a cluster, this field indicates the node for which the statistics are displayed. To display statistics for another node, click the change button.
change	Displays the Select a Cluster Node page in which you can select a different node for which to display statistics.

Field/Item	Description
Last Reset	Displays the date and time the statistics on this page were last reset. To reset the statistics to zero, click the reset button.
Last Refreshed	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Receive Throughput Rate	The receive rate in bytes/second for the Instantaneous (current), Peak, and Total throughput.
Transmit Throughput Rate	The transmit rate in bytes/second for the Instantaneous (current), Peak, and Total throughput.
Total Errors	Lists the number of errors logged on the Fibre Channel ports. The following types of errors are reported: Loss of Signal, Bad Receive Character, Loss of Sync, Link Fail, Receive EOF, Discarded Frames, BAD CRCs, and Protocol Errors.
Congestion	Displays congestion rates. Includes the Instantaneous (current), Peak (over the past 24 hours), and Average (over the past 24 hours) rates.

File and block protocol statistics

The server provides statistics to monitor data access by way of the following network protocols:

- Network File System (NFS)
- Common Internet File System (CIFS)
- File Transfer Protocol (FTP)
- Internet Small Computer System Interface (iSCSI)

Displaying NFS statistics

NFS statistics display activity since the last server reboot or since NFS statistics were last reset. They are updated every 10 seconds.

Procedure

1. Navigate to **Home > Status & Monitoring > NFS Statistics**.

This page displays the current number of RPC calls of different types that clients have issued to the NAS server/cluster node.

File Services
[Home](#) > [File Services](#) > NFS Statistics

NFS Statistics

Cluster Node: Group1-node1
[change...](#)

Last Reset:
2014-05-28 17:39:05 (UTC-0700)
[reset](#)
Last Refreshed: 2014-06-13 13:07:57 (UTC-0700)

	Version 2	Version 3	Version 4
Null	0	49702	0
GetAttr	0	675558	0
SetAttr	0	237	0
Lookup	0	1795432	0
ReadLink	0	0	0
Read	0	3561759	0
Write	0	17192352	0
Create	0	255	0
Remove	0	255	0
Rename	0	192	0
Link	0	0	0
SymLink	0	0	-
MkDir	0	128	-
RmDir	0	128	-
ReadDir	0	0	0
StatFS/FSStat	0	81352	-
MkNod	-	0	-
ReadDirPlus	-	4924	-
FSInfo	-	2	-
PathConf	-	0	-
Commit	-	0	0
Access	-	171112	0
Compound	-	-	0
Close	-	-	0
DelegPurge	-	-	0
DelegReturn	-	-	0
GetFH	-	-	0
Lock	-	-	0
LockT	-	-	0
LockU	-	-	0
LookupP	-	-	0
NVerify	-	-	0
Open	-	-	0
OpenAttr	-	-	0
OpenConfirm	-	-	0
OpenDowngrade	-	-	0
PutFH	-	-	0
PutPubFH	-	-	0
PutRootFH	-	-	0
Renew	-	-	0
RestoreFH	-	-	0
SaveFH	-	-	0
SecInfo	-	-	0
SetClientid	-	-	0
SetClientidConfirm	-	-	0
Verify	-	-	0
ReleaseLockOwner	-	-	0

Field/Item	Description
Cluster Node	For a cluster node, the node is shown in the Cluster Node field and you can change nodes by clicking the change button.

Field/Item	Description
reset	Click to reset the values displayed on this page
Last Refreshed	Displays the date and time the statistics were last reset.
Access	Gets the file security accesses for a file.
Close	Closes a file.
Commit	Commits the cached data on the server to stable storage.
Compound	Compound operations.
Create	Creates a file or symbolic link.
DelegPurge	Purge delegations awaiting recovery.
FSInfo	Gets static file system state information.
FSStat	Gets dynamic file system state information.
GetAttr	Retrieves the attributes of a file or directory.
Link	Creates a hard link to an object.
Lock	Creates a lock.
LockU	Unlocks a file.
Lookup	Looks up a file name in a directory.
LookUpp	Looks up a parent directory.
MkDir	Creates a directory.
MkNod	Creates a special device node (device file or named pipe).
Null	Does nothing, except to make sure the connection is up.
Open	Opens a regular file.
OpenAttr	Opens the named attribute directory.
OpenConfirm	Confirms open.
OpenDowngrade	Reduces open file access.
PathConf	Retrieves POSIX information for the file system.
PutPubFH	Sets public file handle.
PutRootFH	Sets root file handle.
Read	Reads data from a file.
ReadDir	Reads from a directory.
ReadDirPlus	Performs an expanded read from a directory.
ReadLink	Reads the data associated with a symbolic link.
Remove	Removes a file.
Rename	Renames a file or directory.

Field/Item	Description
Renew	Renews a lease.
RestoreFH	Restores saved file handle.
RmDir	Removes a directory.
SaveFH	Saves current file handle.
SecInfo	Obtains available security.
SetAttr	Sets the attributes of a file or directory.
SetClientIdConfirm	Confirms client ID.
StatFS	Gets dynamic file system state information.
SymLink	Creates a symbolic link.
Verify	Verifies same attributes.
Write	Writes data to a file.

- For a cluster node, the node is shown in the **Cluster Node** field, and you can change nodes by clicking the **change** button. These statistics are updated every 10 seconds.
- Click **reset** to reset all the values displayed on this page to zero.

Displaying CIFS statistics

CIFS statistics display SMB activity since the last server reboot or since CIFS statistics were last reset. They are updated every 10 seconds.

Procedure

1. Navigate to **Home > File Services > CIFS Statistics** to display number of current clients and the number of CIFS calls that clients have sent to the server.

File Services [Home](#) > [File Services](#) > CIFS Statistics

CIFS Statistics

Cluster Node: Group1-node1 [change...](#)

Last Refreshed: 2014-06-13 11:45:05 (UTC-0700)

SMB1 statistics

Last Reset: 2014-05-28 17:39:05 (UTC-0700) [reset](#)

Current number of connections: 0
Current number of shares mapped: 0

Protocol Requests

Mkdir: 0	Rmdir: 0	Open: 0
Create: 0	Close: 0	Flush: 0
Unlink: 0	Rename: 0	Getatr: 0
Setatr: 0	Read: 0	Write: 0
Lock: 0	Unlock: 0	CTemp: 0
Mknew: 0	Chkpth: 0	Exit: 0
Lseek: 0	ReadBraw: 0	WriteBraw: 0
SetatrE: 0	GetatrE: 0	LockingX: 0
Trans: 0	Echo: 0	WriteClose: 0
OpenX: 0	ReadX: 0	WriteX: 0
Trans2: 0	FindClose: 0	Tdis: 0
NegProt: 0	SessSetupX: 0	UlogoffX: 0
TconX: 0	Dskattr: 0	Search: 0
NTtrans: 0	NTtrans: 0	NTcreateX: 0
NTcancel: 0	Link: 0	Trans2_open2: 0
Trans2_findFirst2: 0	Trans2_findNext2: 0	Trans2_queryFsInfo: 0
Trans2_queryPathInfo: 0	Trans2_setPathInfo: 0	Trans2_queryFileInfo: 0
Trans2_setFileInfo: 0	Trans2_fsctl: 0	Trans2_ioctl2: 0
Trans2_findNotifyFirst: 0	Trans2_findNotifyNext: 0	Trans2_createDir: 0
Trans2_sessionSetup: 0	Trans2_getDfsReferral: 0	Trans2_reportDfsInconsistency: 0

SMB2 statistics

Last Reset: 2014-05-28 17:37:37 (UTC-0700) [reset](#)

Current number of connections: 0
Current number of shares mapped: 0

Protocol Requests

SMB2_negotiate: 0	SMB2_sessionSetup: 0	SMB2_logoff: 0
SMB2_treeConnect: 0	SMB2_treeDisconnect: 0	SMB2_create: 0
SMB2_close: 0	SMB2_flush: 0	SMB2_read: 0
SMB2_write: 0	SMB2_lock: 0	SMB2_ioctl: 0
SMB2_cancel: 0	SMB2_echo: 0	SMB2_queryDirectory: 0
SMB2_changeNotify: 0	SMB2_queryInfo: 0	SMB2_setInfo: 0
SMB2_oplockBreak: 0		

Protocol Events

SMB2_transportDisconnects: 0	SMB2_unorphanedFileReopens: 0	SMB2_durReopenedFileidAllocFailures: 0
SMB2_durOrphanFileReopenFailures: 0	SMB2_durPreserveOrphanFailures: 0	SMB2_durReconnects: 0

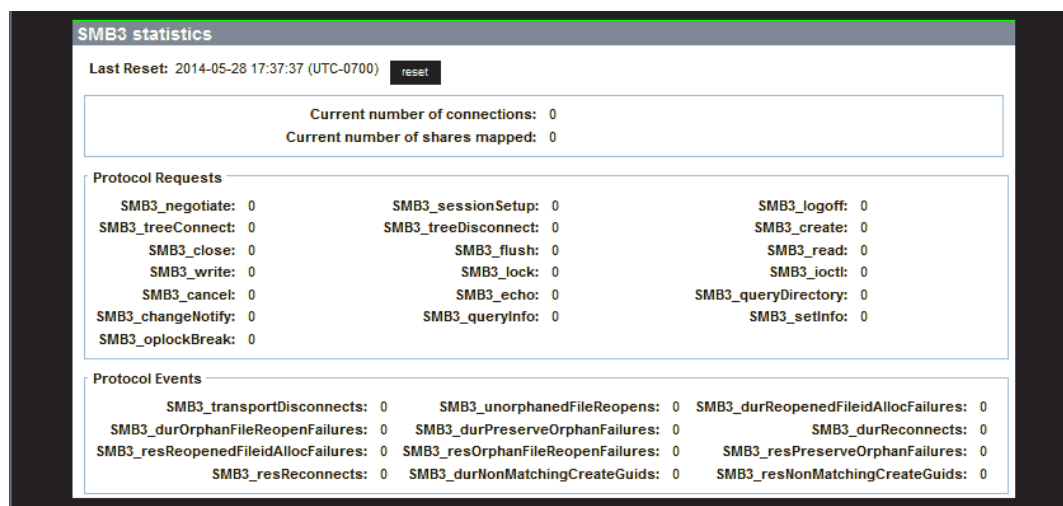
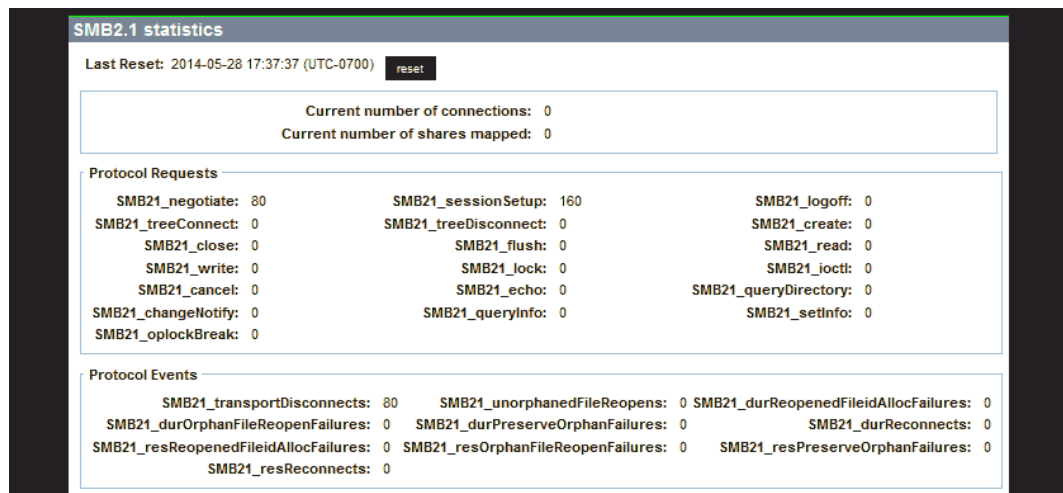


Table 6-1 SMB1 statistics

Field/Item	Description
Cluster Nodes	When connected to a cluster, this field indicates the node for which the statistics are displayed. To display statistics for another node, click the change button.
Last Refreshed	Displays the date, time, and UTC offset for when the statistics on this page were last updated.
Last Reset	Displays the date and time the statistics on this page were last reset.
Time	Displays the UTC offset of the date and time the statistics on this page were last reset.
reset	Resets the statistics to zero for the SMB1 statistics section.

Field/Item	Description
Current number of connections	Displays the current number of SMB1 (CIFS v1) connections.
Current number of shares mapped	Displays the current number of CIFS shares being accessed through the current connections.
Protocol Requests	
Mkdir	Creates a new directory.
Create	Creates a new file or opens an existing one.
Unlink	Deletes a file.
Setattr	Sets the attributes of a file or directory.
Lock	Takes out a byte-range lock on a file.
Mknew	Creates a new file.
Lseek	Sets the file pointer to a given offset in the file.
SetattrE	Sets the attributes of a file or directory.
Trans	Multifunction command for operating subfunctions.
OpenX	Creates a new file or opens an existing one.
Trans2	Multifunction command for operating subfunctions.
NegProt	Negotiates the protocol with which the client and server will communicate.
TconX	Connects the client to a file system resource.
NTtrans	Multifunction command for operating subfunctions.
NTcancel	Cancels an outstanding operation.
Trans2_findFirst2	Begin a search for files.
Trans2_queryPathInfo	Get information about the named file or directory.
Trans2_setFileInfo	Set file information by handle.
Trans2_findNotifyFirst	Commence monitoring changes on a file or directory.
Trans2_sessionSetup	Set up a session with expanded security.
Rmdir	Removes a directory.
Close	Closes a file.
Rename	Renames a file or directory.
Read	Reads data from a file.
Unlock	Releases a byte-range lock on a file.
Chkpth	Checks that the specified directory path exists.
ReadBraw	Reads a block of data with no CIFS header.

Field/Item	Description
GetatrE	Retrieves the expanded attributes of a file or directory.
Echo	Pings the server.
ReadX	Reads data from a file.
FindClose	Closes a CIFS FindFirst subfunction.
SessSetupX	Logs the client in to a CIFS session.
Dskattr	Retrieves file system attributes.
NTtranss	Multifunction command for operating subfunctions.
Link	Creates a hard link to an object.
Trans2_findNext2	Resume a search for files.
Trans2_setPathInfo	Set information about a named file or directory.
Trans2_fsctl	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.
Trans2_findNotifyNext	Continue monitoring changes on a file or directory.
Trans2_GetDfsReferral	Get a DFS referral.
Open	Creates a new file or opens an existing one.
Flush	Instructs the server to flush cached information on a file.
Getattr	Retrieves the attributes of a file or directory.
Write	Writes data to a file.
CTemp	Creates a temporary file with a random server-generated name.
Exit	Used by a process when it exits. Currently unsupported.
WriteBraw	Write a block of data with no CIFS header.
LockingX	Locks or unlocks a range of bytes in a file.
WriteClose	Writes data to a file and then closes the file.
WriteX	Writes data to a file.
Tdis	Breaks a connection that a TconX call previously established.
UlogoffX	Breaks a connection that a SessSetupX call previously established.
Search	Lists the files in a directory.
NTcreateX	Creates a new file or opens an existing one.
Trans2_open2	Create a file that has expanded attributes.
Trans2_queryFsInfo	Get information about a file system.
Trans2_queryFileInfo	Get information about a file handle.
Trans2_ioctl2	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.

Field/Item	Description
Trans2_creatDir	Create a directory that has expanded attributes.
Trans2_reportDfsInconsistency	Report an inconsistency in DFS knowledge.

Table 6-2 SMB2 statistics

Field/Item	Description
Last Reset	Displays the date and time the statistics on this page were last reset.
Time	Displays the UTC offset of the date and time the statistics on this page were last reset.
reset	Resets the statistics to zero for the SMB2 statistics section.
Current number of connections	Displays the current number of SMB2 (CIFS v2) connections.
Current number of shares mapped	Displays the current number of CIFS shares being accessed through the current connections.
Protocol Requests	
SMB2_negotiate	Notify the server what dialects of the SMB 2.0 Protocol the client can process.
SMB2_treeConnect	Request to access to a particular share on the server.
SMB2_close	Close the named resource (pipe or file).
SMB2_write	Write data to the file or named pipe on the server.
SMB2_cancel	Cancels a previously sent message on the same SMB2 transport connection.
SMB2_changeNotify	Change notifications on a directory.
SMB2_oplockBreak	Server notification that the underlying object store indicates that an oplock is being broken, meaning that there is (or will be) a change in the oplock level.
SMB2_sessionSetup	Request for a new authenticated session within a new or existing SMB 2.0 Protocol transport connection to the server.
SMB2_treeDisconnect	Request to terminate the access to the specified tree.
SMB2_flush	Flush all cached file information for a specified open of a file to the persistent store that backs the file.
SMB2_lock	Lock or unlock portions of a file.

Field/Item	Description
SMB2_echo	Determine if a server is processing requests.
SMB2_queryInfo	A request for information on a file, named pipe, or underlying volume.
SMB2_logoff	Terminate the named session.
SMB2_create	Either create a file or access an existing file.
SMB2_read	Request for a read operation on a specified file.
SMB2_ioctl	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.
SMB2_queryDirectory	Get a directory enumeration on an open directory.
SMB2_setInfo	Set information on a file or underlying file system.
Protocol Events	
SMB2_transportDisconnects	The number of times the HNAS node has seen an unexpected network transport disconnect on an SMB2 connection.
SMB2_durOrphanFileReopenFailures	The number of times a client attempted to reopen an orphaned durable file has failed.
SMB2_unorphanedFileReopens	The number of times that a client has tried to reopen a file.
SMB2_durPreserveOrphanFailures	The number of times that a durable file could not be preserved.
SMB2_reopenedFileidAllocationFailures	The number of times that the server has been unable to allocate a new SMB2_FILEID when reopening an orphaned durable file.
SMB2_durReconnects	The number of times that a client has successfully reopened an orphaned durable file.

Table 6-3 SMB2.1 statistics

Field/Item	Description
Last Reset	Displays the date and time the statistics on this page were last reset.
Time	Displays the UTC offset of the date and time the statistics on this page were last reset.
reset	Resets the statistics to zero for the SMB2.1 statistics section.

Field/Item	Description
Current number of connections	Displays the current number of SMB21 (CIFS v2.1) connections.
Current number of shares mapped	Displays the current number of CIFS shares being accessed through the current connections.
Protocol Requests	
SMB21_negotiate	Notify the server what dialects of the SMB 2.1 Protocol the client can process.
SMB21_treeConnect	Request to access to a particular share on the server.
SMB21_close	Close the named resource (pipe or file).
SMB21_write	Write data to the file or named pipe on the server.
SMB21_cancel	Cancels a previously sent message on the same SMB21 transport connection.
SMB21_changeNotify	Change notifications on a directory.
SMB21_oplockBreak	Server notification that the underlying object store indicates that an oplock is being broken, meaning that there is (or will be) a change in the oplock level.
SMB21_sessionSetup	Request for a new authenticated session within a new or existing SMB 2.1 Protocol transport connection to the server.
SMB21_treeDisconnect	Request to terminate the access to the specified tree.
SMB21_flush	Flush all cached file information for a specified open of a file to the persistent store that backs the file.
SMB21_lock	Lock or unlock portions of a file.
SMB21_echo	Determine if a server is processing requests.
SMB21_queryInfo	A request for information on a file, named pipe, or underlying volume.
SMB21_logoff	Terminate the named session.
SMB21_create	Either create a file or access an existing file.
SMB21_read	Request for a read operation on a specified file.
SMB21_ioctl	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.
SMB21_queryDirectory	Get a directory enumeration on an open directory.

Field/Item	Description
SMB21_setInfo	Set information on a file or underlying file system.
Protocol Events	
SMB21_transportDisconnects	The number of times the HNAS node has seen an unexpected network transport disconnect on an SMB2 connection.
SMB21_durOrphanFileReopenFailures	The number of times a client attempted to reopen an orphaned durable file has failed.
SMB21_resReopenedFileidAllocFailures	The number of times that the server has been unable to allocate a new SMB21_FILEID when reopening an orphaned durable file.
SMB21_resReconnects	The number of times that a client has successfully reopened an orphaned durable file.
SMB21_unorphanedFileReopens	The number of times that a client has tried to reopen a file.
SMB21_durPreserveOrphanFailures	The number of times that a durable file could not be preserved.
SMB21_resOrphanFileReopenFailures	The number of times that an orphaned file could not be reopened.
SMB21_durReopenedFileidAllocFailures	The number of times that an orphaned durable file id could not be re-allocated.
SMB21_durReconnects	The number of times that a client has successfully reopened an orphaned durable file.
SMB21_resPreserveOrphanFailures	The number of times that a resilient file could not be preserved.

Table 6-4 SMB3 statistics

Field/Item	Description
Last Reset	Displays the date and time the statistics on this page were last reset.
Time	Displays the UTC offset of the date and time the statistics on this page were last reset.
reset	Resets the statistics to zero for the SMB3 statistics section.
Current number of connections	Displays the current number of SMB3 (CIFS v3.0) connections.
Protocol Requests	

Field/Item	Description
Current number of shares mapped	Displays the current number of CIFS shares being accessed through the current connections.
SMB3_negotiate	Notify the server what dialects of the SMB 3.0 Protocol the client can process.
SMB3_treeConnect	Request to access to a particular share on the server.
SMB3_close	Close the named resource (pipe or file).
SMB3_write	Write data to the file or named pipe on the server.
SMB3_cancel	Cancels a previously sent message on the same SMB3 transport connection.
SMB3_changeNotify	Change notifications on a directory.
SMB3_oplockBreak	Server notification that the underlying object store indicates that an oplock is being broken, meaning that there is (or will be) a change in the oplock level.
SMB3_sessionSetup	Request for a new authenticated session within a new or existing SMB 3.0 Protocol transport connection to the server.
SMB3_treeDisconnect	Request to terminate the access to the specified tree.
SMB3_flush	Flush all cached file information for a specified open of a file to the persistent store that backs the file.
SMB3_lock	Lock or unlock portions of a file.
SMB3_echo	Determine if a server is processing requests.
SMB3_queryInfo	A request for information on a file, named pipe, or underlying volume.
SMB3_logoff	Terminate the named session.
SMB3_create	Either create a file or access an existing file.
SMB3_read	Request for a read operation on a specified file.
SMB3_ioctl	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.
SMB3_queryDirectory	Get a directory enumeration on an open directory.

Field/Item	Description
SMB3_setInfo	Set information on a file or underlying file system.
Protocol Events	
SMB3_transportDisconnects	The number of times the HNAS node has seen an unexpected network transport disconnect on an SMB3 connection.
SMB3_durOrphanFileReopenFailures	The number of times a client attempted to reopen an orphaned durable file has failed.
SMB3_resReopenedFileidAllocFailures	The number of times that the server has been unable to allocate a new SMB3_FILEID when reopening an orphaned durable file.
SMB3_resReconnects	The number of times that a client has successfully reopened an orphaned durable file.
SMB3_unorphanedFileReopens	The number of times that a client has tried to reopen a file.
SMB3_durPreserveOrphanFailures	The number of times that a durable file could not be preserved.
SMB3_resOrphanFileReopenFailures	The number of times that an orphaned file could not be reopened.
SMB3_durNonMatchingCreateGuids	The number of times that a client presented an incorrect create guid for a durable file when reconnecting.
SMB3_durReopenedFileidAllocFailures	The number of times that an orphaned durable file id could not be re-allocated.
SMB3_durReconnects	The number of times that a client has successfully reopened an orphaned durable file.
SMB3_resPreserveOrphanFailures	The number of times that a resilient file could not be preserved.
SMB3_resNonMatchingCreateGuids	The number of times that a client presented an incorrect create guid for a resilient file when reconnecting.

2. Click **reset** in either the SMB1 or SMB2 section to set all values in that section to zero.

Displaying FTP statistics

FTP statistics display activity since the last server reboot or since FTP statistics were last reset. They are updated every 10 seconds.

Procedure

1. Navigate to **Home > Status & Monitoring > FTP Statistics** to display the **FTP Statistics** page.

[File Services](#) > [Home](#) > [File Services](#) > FTP Statistics

FTP Statistics

Cluster Node: Group1-node1
[change...](#)

Last Reset: 2014-05-28 17:39:05 (UTC-0700) [reset](#) Last Refreshed: 2014-06-13 13:09:28 (UTC-0700)

Sessions
Current Active Sessions: 0
Total Sessions: 0
Current Active Transfers: 0

Commands
Commands Issued from Clients: 0
Total Replies Sent to Clients: 0
Total Bytes Received in Commands: 0
Total Bytes Sent in Replies: 0

Files
Files Incoming for Active Sessions: 0
Total Files Incoming: 0
Files Outgoing for Active Sessions: 0
Total Files Outgoing: 0

Data Bytes
Data Bytes Incoming for Active Sessions: 0
Total Data Bytes Incoming: 0
Data Bytes Outgoing for Active Sessions: 0
Total Data Bytes Outgoing: 0

Field/Item	Description
Cluster Node	For a cluster node, the node is shown in the Cluster Node field and you can change nodes by clicking the change button.
Reset	You can reset the values displayed on this page to zero by clicking reset . All values are reset except the ones concerning active sessions, that is, number of active sessions, number of files incoming/outgoing for active sessions and number of bytes incoming/outgoing for active sessions.
Last Refreshed	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Sessions	
Current Active Sessions	Currently active FTP sessions.
Total Sessions	Total FTP sessions since last server restart or statistics reset.
Current Active Transfers	Currently active FTP transfers.
Commands	
Commands Issued from Clients	Number of commands sent by clients.
Total Replies Sent to Clients	Number of replies sent to clients.

Field/Item	Description
Total Bytes Received in Commands	Bytes in commands that clients have sent to the FTP server.
Total Bytes Sent in Replies	Bytes in replies that the FTP server has sent to clients.
Files	
Bytes in replies that the FTP server has sent to clients.	Files that clients have transferred to the FTP server in currently active sessions.
Total Files Incoming	Files that clients have transferred to the FTP server since last server restart or statistics reset.
File Outgoing for Active Sessions	Files that the FTP server has transferred to clients in currently active sessions.
Total Files Outgoing	Files that the FTP server has transferred to clients since last server restart or statistics reset.
Data bytes	
Data Bytes Incoming for Active Sessions	Bytes of data that clients have transferred to the FTP server in currently active sessions.
Total Data Bytes Incoming	Bytes of data that clients have transferred to the server since last server restart or statistics reset.
Data Bytes Outgoing for Active Sessions	Bytes of data that the FTP server has transferred to clients in currently active sessions.
Total Data Bytes Outgoing	Bytes of data that the server has transferred to clients since last server restart or statistics reset.

2. Click **reset** to set the values displayed on this page to zero. All values are reset except the ones concerning active sessions; for example, the number of active sessions, number of files incoming/outgoing for active sessions and number of bytes incoming/outgoing for active sessions.

Displaying iSCSI statistics

The **iSCSI Statistics** page provides a summary of the iSCSI and SCSI requests on a NAS server/cluster node. These statistics are updated every 10 seconds.

Procedure

1. Navigate to **Home > File Services > iSCSI Statistics** to display the **iSCSI Statistics** page.

[File Services](#) [Home](#) > [File Services](#) > iSCSI Statistics

iSCSI Statistics

Cluster Node: Group1-node1
[change...](#)

Last Reset: 2014-05-29 00:14:12 (UTC-0700) [reset](#) Last Refreshed: 2014-06-13 13:08:51 (UTC-0700)

Current Sessions: 0
Current Connections: 0

iSCSI Requests

NopOut:	0
SCSICommand:	0
TaskManagement:	0
Login:	0
Text:	0
SCSIDataOut:	0
Logout:	0

SCSI Requests

TEST UNIT READY:	0
REQUEST SENSE:	0
FORMAT UNIT:	0
READ(6):	0
WRITE(6):	0
INQUIRY:	0
MODE SELECT(6):	0
RESERVE(6):	0
RELEASE(6):	0
MODE SENSE(6):	0
START STOP UNIT:	0
READ CAPACITY(10):	0
READ(10):	0
WRITE(10):	0
WRITE AND VERIFY(10):	0
VERIFY(10):	0
SYNCHRONIZE CACHE(10):	0
MODE SELECT(10):	0
RESERVE(10):	0
RELEASE(10):	0
PERSISTENT RESERVE IN:	0
PERSISTENT RESERVE OUT:	0
READ(16):	0
WRITE(16):	0
VERIFY(16):	0
SERVICE ACTION IN(16):	0
REPORT LUNS:	0

Field/Item	Description
Cluster Node	For a cluster node, the node is shown in the Cluster Node field and you can change nodes by clicking the change button.
reset	Click to reset the values displayed on this page
Last Refreshed	Displays the date and time the statistics were last reset.
Current Connections	The current number of iSCSI connections to the server.
Current Number of Session	The number of iSCSI sessions currently hosted by the server.
iSCSI requests	
NopOut	No operation.
Task Management	Requests used for task management functions.
Text	Requests used to negotiate behavior.
Logout	Logout requests.
SCSICommand	Carries a SCSI command.
Login	Login requests.
SCSIDataOut	Requests containing SCSI data.
iSCSI requests	
TestUnitReady	Tests that the target is ready to receive commands.
Read(6)	Reads data.
ModeSelect(6)	Configure SCSI behavior.
Release(6)	Releases (unlocks) a LU reservation.
StartStopUnit	Warm reboots the target.
Read(10)	Reads data.
Verify(10)	Verifies data.
ModeSelect(10)	Configure SCSI behavior.
Release(10)	Releases (unlocks) a LU reservation.
RequestSense	Requests state information.
Inquiry	Requests device information.
Reserve(6)	Reserves (locks) a LU for exclusive access.
ModeSense(6)	Requests SCSI configuration information.
ReadCapacity	Reads the size of LU.
Write(10)	Write data.
SynchronizeCache	Flushes cached data to disk.
Reserve(10)	Reserves (locks) a LU for exclusive access.

Field/Item	Description
Persistent Reserve In	An inbound reserve that persists even after system reset
Persistent Reserve Out	An outbound reserve that persists even after system reset
ReportLuns	Retrieves a list of available LUs.
Format Unit	Formats a LU.
Write(6)	Writes data.
WriteAndVerify(10)	Writes then verifies data.
Read(16)	Reads data.
Write(16)	Writes data.
Verify(16)	Verifies data.
ServiceActionIn(16)	Performs an extended SCSI command, such as ReadCapacity(16).

- For a cluster node, the node is shown in the Cluster Node field and you can change nodes by clicking the **change** button.
- Click **reset** to set the values displayed on this page to zero.

Data access and performance statistics

The server provides measures and tools for monitoring the impact of network clients on internal resources. In particular, the server provides:

- Server and file system load statistics
- File system NVRAM usage statistics

Server and file system load statistics

In addition to Ethernet and Fibre Channel throughput statistics, server performance can also be measured in operations per second (ops/sec). The Web Manager provides a graphic representation of ops/sec, at two levels:

- Total operations per server.
- Total operations per individual file system.

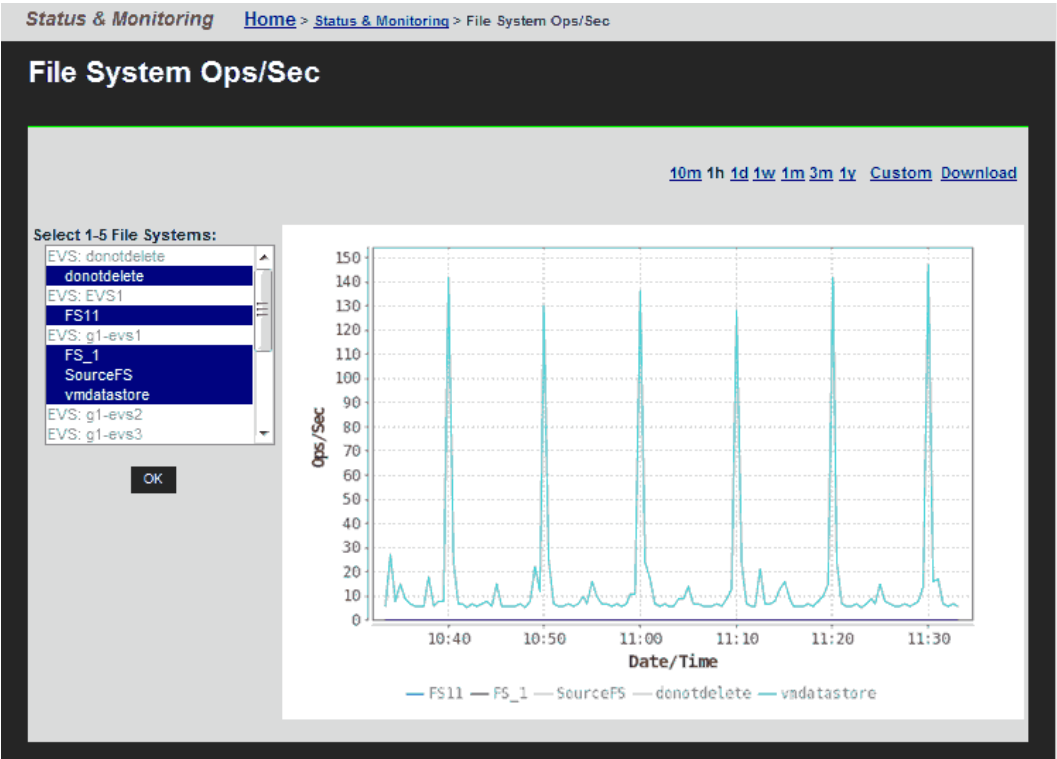
The *total operations on a server* is an aggregate of the operations performed by all file systems hosted by that server.

Understanding the performance profile of servers and individual file systems is especially useful in environments where more than one server is installed, as it enables intelligent relocation of EVSs or file systems to more equally distribute the load among the available servers.

Displaying operations per second (ops/sec) statistics

Procedure

1. Navigate to **Home > Status & Monitoring**, and select either **File System Ops/Sec** or **Node Ops/Sec**.



Item	Description
10m 1h 1d 1w 1m 3m 1y	Click the corresponding link to display graphs for the last 10 minutes (10m), 1 hour (1h), 1 day (1d), 1 week (1w), 1 month (1m), 3 months (3m), or 1 year (1y).
Custom	Click to specify a custom date range, which includes the from and to dates and times. Click OK to initiate the graph.
Download	Download statistics on the graph as a single CSV table.
Select 1-5 File Systems	Displays all the file systems in the currently managed server or cluster. To select file systems, and display the performance data for that file system: <ol style="list-style-type: none">1. On your keyboard, press and hold the Control (Ctrl) key.2. Highlight up to five file systems.3. Click OK.
Ops/Sec	The number of recorded operations per second.

Item	Description
Date/Time	The currently selected date/time range.

Displaying file system NVRAM statistics

The **File System NVRAM Statistics** page displays NVRAM activity.



Note: When an EVS has a Read Cache file system, no NVRAM statistics are presented.

Procedure

1. Navigate to **Home > Storage Management > File System NVRAM Statistics**.

Storage Management [Home](#) > [Storage Management](#) > [NVRAM Statistics](#)

NVRAM Statistics

Cluster Node: Group1-node1

Last Refreshed: 2014-06-13 11:38:28 (UTC-0700)

NVRAM size: 1005 MB
Maximum used: 176 MB
Currently in use: 88 MB

Field/Item	Description
Cluster Node	When connected to a cluster, this field indicates the node for which NVRAM statistics are displayed. To display statistics for another node, click the change button.
change	Displays the Select a Cluster Node page in which you can select a different node for which to display statistics.
Last Refreshed	Displays the date and time this page was refreshed.
NVRAM size	Size of NVRAM buffer, used to preserve data for disk-modifying operations until written to disk. The default is 2 GB.
Maximum used	Maximum amount of the NVRAM buffer that has been used since the node was last started.
Currently in use	Currently in use

Management statistics

A NAS Platform 3080 or NAS Platform 3090/cluster provides the following management statistics:

- Access management statistics for SSC, SNMP, HTTPS, and VSS

- Virus scanning statistics

A NAS Platform Series 2000 or NAS Platform Series 3000/cluster also provides the following management statistics:

- Access management statistics for Telnet, SSC, SSH, SNMP, HTTPS, and VSS
- Virus scanning statistics

Displaying access management statistics

Procedure

1. Navigate to **Home > Status & Monitoring**, and select one of the items in the **Management Access Statistics** section.

Status & Monitoring Home > Status & Monitoring > SNMP Management Statistics			
SNMP Management Statistics			
Cluster Node: Group1-node1			
change...			
Last Reset: 2014-05-28 17:37:34 (UTC-0700) reset Last Refreshed: 2014-06-09 16:20:56 (UTC-0700)			
	Input	Output	Drops
Packets	0	0	-
Bad Versions	0	-	-
Bad Community Names	0	-	-
Bad Community Uses	0	-	-
Too Bigs	0	0	-
No Such Names	0	0	-
Bad Values	0	0	-
Read Onlys	0	-	-
General Errors	0	0	-
Total Request Varbinds	0	-	-
Total Set Varbinds	0	-	-
Get Requests	0	0	-
Get Nexts	0	0	-
Set Requests	0	0	-
Get Responses	0	0	-
Traps	0	343	-
ASN Parse Errors	0	-	-
Silent Drops	-	-	0
Proxy Drops	-	-	0

Field/Item	Description
Cluster Node	When the server is part of a cluster, this field identifies the node. Click change to change nodes, and view statistics for that node.
Last Reset	The date and time the statistics on this page were reset. Click reset to reset the statistics to zero.
Last Refreshed	The date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Status	The status will be one of the following: <ul style="list-style-type: none">• 'Listening on port x ' - HTTPS or any other management service such as HTTP, SSC, or VSS is enabled and available for service at a configured port.• 'Server disabled' - HTTPS or other management service is disabled.• 'Listening on EVS IP addresses' -The service is running and it is listening on EVS IP addresses. This status does not apply to HTTP, SSC, or VSS.

Field/Item	Description
Maximum Simultaneous Connections	The maximum number of connections at any given time.
Connections	Number of sessions that are currently in progress.
Current Active Sessions	The number of HTTPS sessions that are currently in progress.
Max Sessions	The peak number of concurrent HTTPS sessions.
Total Sessions	The total number of HTTPS sessions.
Rejected Sessions	The number of rejected sessions.
Successful Logins	The total number of successful logins
Failed Logins	The total number of failed login attempts.
Frames	
Frames Transmitted	The total number of frames that the system has sent to clients over an HTTPS connection.
Frames Received	The total number of frames that clients have sent to the system over an HTTPS connection.
Bytes	
Bytes Transmitted	The number of data bytes that the system has sent to clients over an HTTPS connection.
Bytes Received	The number of data bytes that clients have sent to the system over an HTTPS connection.
Current Connections (for each active connection)	
Address	IP address of the connected client.
Age	Time, in seconds, the connection has been active.
Rx Age (s)	Time, in seconds, since the server last received a packet from the client.
Tx Age (s)	Time, in seconds, since the server last sent a packet to the client.
Rx Packets	Time, in seconds, since the server last received a packet from the client.
Tx Packets	Time, in seconds, since the server last sent a packet to the client.

- When the server is part of a cluster, the Cluster Node field identifies the node, and the **change** button allows you to change nodes, and display statistics for that node.
- Click **reset** to set the values displayed on this page to zero.

Displaying SNMP management statistics

The SNMP management statistics page displays the SNMP statistics for the server since the server was last reset. It displays statistics regarding Input, Output, and Drops. These statistics are updated every 10 seconds.

Procedure

1. Navigate to **Home > Status & Monitoring > SNMP Management Statistics**.

Status & Monitoring [Home](#) > [Status & Monitoring](#) > [SSC Management Statistics](#)

SSC Management Statistics

Cluster Node: Group1-node1

Last Reset: 2014-05-28 17:39:05 (UTC-0700) **Last Refreshed:** 2014-06-13 13:43:35 (UTC-0700)

Status: Listening on port 206
Maximum Simultaneous Connections: 5

Connections	Frames	Data Bytes
Current Active Connections: 0	Frames Transmitted: 1441000	Bytes Transmitted: 93558781
Max Connections: 4	Frames Received: 371428	Bytes Received: 17656928
Total Connections: 92745		
Rejected Connections: 0		
Successful Logins: 92745		
Failed Logins: 0		

Current Connections

Address	Age	Rx Age (s)	Tx Age (s)	Rx packets	Tx packets
---------	-----	------------	------------	------------	------------

Field/Item	Description
Cluster Node	For a cluster node, the node is shown in the Cluster Node field and you can change nodes by clicking the change button.
reset	Click to reset the values displayed on this page
Last Refreshed	Displays the date and time the statistics were last reset.
Input	
Packets	SNMP packets the agent has received.
Bad Versions	Packets received that were for an unsupported SNMP version.
Bad Community Names	SNMP messages received using an unknown community name.
Bad Community Uses	SNMP messages received that used an unknown community name.
Too Bigs	Protocol Data Units (PDUs) received containing an error-status field value of tooBig.
No Such Names	PDUs received that contained an error-status field value of nosuchName.
Bad Values	PDUs received containing an error-status field value of badValue.
Read Onlys	PDUs received that contained an error-status field value of ReadOnly. This value is used to detect incorrect SNMP implementations.

Field/Item	Description
General Errors	PDUs received containing an error-status field value of genErr.
Total Request Varbinds	MIB objects successfully retrieved because of valid SNMP Get-Request and Get-Next PDUs.
Total Set Varbinds	MIB objects successfully altered because of valid SNMP Set-Request PDUs.
Get Requests	Get-Request PDUs sent.
Get Nexts	Get-Next PDUs received and processed.
Set Requests	Set-Request PDUs received and processed.
Get Responses	Get-Response PDUs received and processed.
Traps	Trap PDUs received and processed.
ASN Parse Errors	Abstract Syntax Notation (ASN) errors found in SNMP messages received.
Output	
Packets	SNMP packets the agent has sent.
Too Bigs	Sent PDUs receiving an error-status field value of tooBig.
No Such Names	Sent PDUs receiving an error-status field value of noSuchName.
Bad Values	Sent PDUs receiving an error-status field value of badValue.
General Errors	Sent PDUs receiving an error-status field value of genErr.
Get Requests	Get-Request PDUs sent.
Get Nexts	Get-Next PDUs sent.
Set Requests	Set-Request PDUs sent.
Get Responses	Get-Response PDUs sent.
Traps	Trap PDUs sent.
Drops	
Silent Drops	PDUs delivered but silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests.
Proxy Drops	PDUs delivered but silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned.

2. When the server is part of a cluster, the Cluster Node field identifies the node, and the **change** button allows you to change nodes, and display statistics for that node.
3. Click **reset** to set the values displayed on this page to zero.

Displaying HTTPS management statistics

The **HTTPS Management Statistics** page displays the HTTPS statistics for the server since the server was last reset. It displays statistics regarding sessions, data sent/received, and connections. These statistics are updated every 10 seconds.

Procedure

1. Navigate to **Home > Status & Monitoring > HTTPS Management Statistics**.

Status & Monitoring [Home](#) > [Status & Monitoring](#) > [HTTPS Management Statistics](#)

HTTPS Management Statistics

Cluster Node: Group1-node1
[change...](#)

Last Reset: 2014-05-28 17:39:13 (UTC-0700) [reset](#) Last Refreshed: 2014-06-13 13:42:34 (UTC-0700)

Status: Listening on port 8443
Maximum Simultaneous Connections: 50

Connections	Frames	Data Bytes
Current Active Connections: 3	Frames Transmitted: 6508978	Bytes Transmitted: 32247406995
Max Connections: 8	Frames Received: 4278492	Bytes Received: 3242490491
Total Connections: 10731		
Rejected Connections: 0		
Successful Logins: 10731		
Failed Logins: 0		

Current Connections					
Address	Age	Rx Age (s)	Tx Age (s)	Rx packets	Tx packets
192.0.2.1	6 hours 42 minutes 36 seconds	13	24156	26917	20194
192.0.2.1	4 hours 42 minutes 35 seconds	148	16955	19628	14679
192.0.2.1	42 minutes 34 seconds	0	2554	3412	2546

Field/Item	Description
Cluster Node	When the server is part of a cluster, this field identifies the node. Click change to change nodes, and view statistics for that node.
Last Reset	The date and time the statistics on this page were reset. Click reset to reset the statistics to zero.
Last Refreshed	The date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Status	<p>The status will be one of the following:</p> <ul style="list-style-type: none"> 'Listening on port x ' - HTTPS or any other management service such as HTTP, SSC, or VSS is enabled and available for service at a configured port. 'Server disabled' - HTTPS or other management service is disabled. 'Listening on EVS IP addresses' -The service is running and it is listening on EVS IP addresses. This status does not apply to HTTP, SSC, or VSS.
Maximum Simultaneous Connections	The maximum number of connections at any given time.
Connections	
Current Active Sessions	The number of HTTPS sessions that are currently in progress.

Field/Item	Description
Max Sessions	The peak number of concurrent HTTPS sessions.
Total Sessions	The total number of HTTPS sessions.
Rejected Sessions	The total number of failed attempts to establish an HTTPS connection. A connection might fail because the client does not have the required permissions or because the maximum number of concurrent sessions are already in progress.
Successful Logins	The total number of successful logins
Failed Logins	The total number of failed login attempts.
Frames	
Frames Transmitted	The total number of frames that the system has sent to clients over an HTTPS connection.
Frames Received	The total number of frames that clients have sent to the system over an HTTPS connection.
Data Bytes	
Bytes Transmitted	The number of data bytes that the system has sent to clients over an HTTPS connection.
Bytes Received	The number of data bytes that clients have sent to the system over an HTTPS connection.
Current Connections (for each active connection)	
Address	IP address of the connected client.
Age	Time, in seconds, the connection has been active.
Rx Age (s)	Time, in seconds, since the server last received a packet from the client.
Tx Age (s)	Time, in seconds, since the server last sent a packet to the client.
Rx Packets	Time, in seconds, since the server last received a packet from the client.
Tx Packets	Time, in seconds, since the server last sent a packet to the client.

2. When the server is part of a cluster, the Cluster Node field identifies the node, and the **change** button allows you to change nodes, and display statistics for that node.
3. Click **reset** to set the values displayed on this page to zero.

Displaying VSS management statistics

The **VSS Management Statistics** page displays the VSS statistics for the server since the server was last reset. It displays statistics regarding sessions, frames sent/received, and data sent/received. These statistics are updated every 10 seconds.

Procedure

1. Navigate to **Home > Status & Monitoring > VSS Management Statistics**.

[Status & Monitoring](#) [Home](#) > [Status & Monitoring](#) > VSS Management Statistics

VSS Management Statistics

Cluster Node: Group1-node1
[change...](#)

Last Reset: 2014-05-28 17:39:05 (UTC-0700) [reset](#) Last Refreshed: 2014-06-13 13:43:09 (UTC-0700)

Status: Listening on port 202
Maximum Simultaneous Connections: 5
Activity: No Activity

Current Connections					
Address	Age	Rx Age (s)	Tx Age (s)	Rx packets	Tx packets

Field/Item	Description
Sessions	
Current Active Sessions	The number of HTTPS sessions that are currently in progress.
Max Sessions	The peak number of concurrent HTTPS sessions.
Total Sessions	The total number of HTTPS sessions.
Rejected Sessions	The total number of failed attempts to establish an HTTPS connection. A connection might fail because the client does not have the required permissions or because the maximum number of concurrent sessions are already in progress.
Frames	
Frames Transmitted	The total number of frames that the system has sent to clients over an HTTPS connection.
Frames Received	The total number of frames that clients have sent to the system over an HTTPS connection.
Data Bytes	
Bytes Transmitted	The number of data bytes that the system has sent to clients over an HTTPS connection.
Bytes Received	The number of data bytes that clients have sent to the system over an HTTPS connection.

2. When the server is part of a cluster, the Cluster Node field identifies the node, and the **change** button allows you to change nodes, and display statistics for that node.

3. Click **reset** to set the values displayed on this page to zero.

Displaying virus scanning statistics

The **Virus Statistics** page summarizes virus scanning activity.



Note: Files will only be deleted, repaired, or quarantined if the virus scan engine has been configured to do so.

Procedure

1. Navigate to **Home > Data Protection > Virus Statistics**.

[Data Protection](#) [Home](#) > [Data Protection](#) > Virus Statistics

Virus Statistics

EVS: g1-evs3 [change...](#)

Last Reset: 2013-07-18 12:15:21 (UTC-0700) [reset](#) Last Refreshed: 2014-06-09 16:23:42 (UTC-0700)

Statistics

Number of virus scans:	0
Number of clean scans:	0
Number of errored scans:	0

Additional statistics

Number of infections found:	0
Action taken:	
Number of infections repaired:	0
Number of files deleted:	0
Number of files quarantined:	0



Note: When a virus is detected, a severe event is placed in the **Event Log**, identifying the path of the infected file and the IP address of the client machine that wrote the file.



Important: Files will only be deleted, repaired or quarantined if the virus scan engine has been configured to do so.

Field/Item	Description
EVS	Identifies the currently selected EVS; click change to display statistics for a different EVS.
Last Reset	Displays the last time the statistics were reset to zero. Click the reset button to reset all values to zero.
Last Updated	Displays the date and time the statistics were last updated
Number of virus scans	Number of times files have been scanned for viruses.
Number of clean scans	Number of times files have been scanned with no viruses detected.

Field/Item	Description
Number of errored scans	Number of times a failure occurred while scanning a file.
Additional statistics (not supported on every virus scan engine)	
Number of infections found	Number of times files have been scanned and detected infections are found.
Number of infections repaired	Number of times the virus scan engine has been able to repair infections found.
Number of files deleted	Number of deleted files because they contain irreparable infections.
Number of files quarantined	Number of files quarantined because they contain irreparable infections.

2. When the server is part of a cluster, the Cluster Node field identifies the node, and the **change** button allows you to change nodes, and display statistics for that node.

Event logging and notification

The server provides a comprehensive event logging and alert mechanism and auxiliary devices in the storage subsystem automatically direct any events and SNMP traps to the server (or can be configured to do so).

All event messages generated by the server (including those issued by its auxiliary devices) are logged into an event log, which can be downloaded and cleared by the system administrator. The event log provides a record of past events that have occurred on the server, for use in trend/fault analysis.

Event message severity can be changed, and messages can be suppressed entirely, using the `event-log-filter` command. Using `event-log-filter`, you can specify that a command is to be run whenever a specified message is logged. For more information on the `event-log-filter` command, enter `man event-log-filter`, or refer to the *Command Line Reference*.

The server can also be configured for automated notification according to predefined severity categories, including daily summary and status notification. With automated notification enabled, the system will notify selected personnel when an event is generated, based on the level of severity of the event. 24x7 automated notifications allow Hitachi Data Systems Support Center personnel to proactively monitor the health of the system and address any issues that may arise.

Using the event log

The server continuously monitors temperature, fans, power supply units, and disk drives in the cabinet. Each time an event occurs (for example, a disk failure or a possible breach of security, the system records it in an event log). The event log can be displayed, filtered, and saved as a permanent record.

The log can contain a maximum of 10,000 events. Once the event log limit has been reached, each new event replaces the oldest event in the log.

Displaying and filtering the event log

Procedure

1. Navigate to **Home > Status & Monitoring > Event Log** to display the **Event Log Management** page.

Status & Monitoring [Home](#) > [Status & Monitoring](#) > [Event Log](#)

Event Log

filter Current Time: 2014-07-25 10:25:23 (UTC-0700) events 1-20 of 10000 : Page: 1 2 3 Show 20 items per page

ID	Severity	Cluster Node	Date/Time	Event
5167	Warning	Group1-node1	2014-07-25 09:37:16 (UTC-0700)	No EVS is running on cluster node 1.
8477	Warning	Group1-node2	2014-07-25 08:00:55 (UTC-0700)	[group1-smu] SMU cannot reach the RADIUS server 'R-Server03'.
8477	Warning	Group1-node2	2014-07-25 08:00:55 (UTC-0700)	[group1-smu] SMU cannot reach the RADIUS server 'RADIUS02'.
8477	Warning	Group1-node2	2014-07-25 08:00:55 (UTC-0700)	[group1-smu] SMU cannot reach the RADIUS server 'RadServ01'.
7162	Information	Group1-node2	2014-07-25 06:40:24 (UTC-0700)	Snapshot deleted ("Dedupe incremental indexing" on nolan).
5903	Information	Group1-node2	2014-07-25 06:40:24 (UTC-0700)	Dedupe: Finished daily incremental job on file system 'nolan'. 16 B indexed, 0 B duplicates found (0%).
7161	Information	Group1-node2	2014-07-25 06:40:24 (UTC-0700)	Snapshot marked for deletion ("Dedupe incremental indexing" on nolan).
7160	Information	Group1-node2	2014-07-25 06:40:18 (UTC-0700)	Snapshot created ("Dedupe incremental indexing" on nolan).
7162	Information	Group1-node2	2014-07-25 04:40:40 (UTC-0700)	Snapshot deleted ("Dedupe incremental indexing" on FS_1).
5903	Information	Group1-node2	2014-07-25 04:40:40 (UTC-0700)	Dedupe: Finished daily incremental job on file system 'FS_1'. 16 B indexed, 0 B duplicates found (0%).
7161	Information	Group1-node2	2014-07-25 04:40:40 (UTC-0700)	Snapshot marked for deletion ("Dedupe incremental indexing" on FS_1).
7160	Information	Group1-node2	2014-07-25 04:40:22 (UTC-0700)	Snapshot created ("Dedupe incremental indexing" on FS_1).
8221	Warning	Group1-node2	2014-07-25 00:19:49 (UTC-0700)	Chassis disk '/dev/sda' Smartd message type 'SelfTest' message text 'Device: /dev/sda, new Self-Test Log error at hour timestamp 29975'.
7162	Information	Group1-node2	2014-07-25 00:00:19 (UTC-0700)	Snapshot deleted ("2014-07-13_0000-0700.daily" on vmdatstore).
7161	Information	Group1-node2	2014-07-25 00:00:18 (UTC-0700)	Snapshot marked for deletion ("2014-07-13_0000-0700.daily" on vmdatstore).
7160	Information	Group1-node2	2014-07-25 00:00:18 (UTC-0700)	Snapshot created ("2014-07-25_0000-0700.daily" on vmdatstore).
5434	Information	Group1-node2	2014-07-25 00:00:05 (UTC-0700)	NDMP(4): Starting operation for client 172.31.60.40, standard NDMP user.
5434	Information	Group1-node2	2014-07-25 00:00:05 (UTC-0700)	NDMP(3): Starting operation for client 172.31.60.40, standard NDMP user.
7162	Information	Group1-node2	2014-07-24 16:40:46 (UTC-0700)	Snapshot deleted ("Dedupe incremental indexing" on protectedsiteA).
5903	Information	Group1-node2	2014-07-24 16:40:46 (UTC-0700)	Dedupe: Finished daily incremental job on file system 'protectedsiteA'. 16 B indexed, 0 B duplicates found (0%).

events 1-20 of 10000 : Page: 1 2 3

Actions: refresh cache download clear all

Shortcuts: [Active Tasks](#)

Look up [SCSI Error Codes](#) (Sense Key/ASC/ASCQ)

2. Click **filter** to open the **Filter** dialog.

The screenshot shows a 'Filter' dialog box with the following fields and options:

- Cluster Node:** A dropdown menu currently set to 'All Cluster Nodes'.
- Event Category:** A dropdown menu currently set to 'All'.
- Event ID:** An empty text input field.
- Event Description:** An empty text input field.
- Severity Levels:** Three checkboxes labeled 'Information', 'Warning', and 'Severe', all of which are checked.
- Time Range:** Two sections, 'From' and 'To', each containing a date and time selection. Both are currently set to '2014-07-25'. The 'From' time is '00:00:00' and the 'To' time is '23:59:59'.
- Buttons:** 'OK' and 'reset' buttons at the bottom.

- In a cluster, specify the cluster node for which to display the log. In the Cluster Node field, you can select the specific node or **All Cluster Nodes**.
- In the Event Category field, select the type of events to be included in the log: **All events**, **System events**, or **Security events**.
System events are events that the system components have logged, such as the failure of a drive. Security events track changes to the security system and identify possible breaches of security.
- You can specify an Event ID that you want included in the log.
- You can specify an Event Description that you want included in the log.
- Select the severity level of the events you want included in the log by filling one or more of the boxes: **Information**, **Warning**, or **Severe**.
- Click **OK** to filter the log events being displayed according to the filter criteria you specified.

3. Click an event to display the cause and resolution.

The screenshot shows the 'Event Log' page in a web interface. At the top, there's a breadcrumb trail: 'Status & Monitoring' > 'Home' > 'Status & Monitoring' > 'Event Log'. Below this is a table with columns: ID, Severity, Cluster Node, Date/Time, and Event. The table lists several events, including one with ID 5513, Severity 'Information', Cluster Node 'Group1-node2', and Date/Time '2014-06-09 15:28:17 (UTC-0700)'. A popup window titled 'Cause And Resolution' is overlaid on the table, showing details for event ID 5513. The popup contains the following text:

Cause: The management network link used for heart beating has failed.

Resolution: Investigate what caused the link to fail, or re-establish connectivity. Contact your support provider if necessary.

4. Click **refresh cache** to clear the SMU's cache, and then repopulate the cache with the relevant objects.
Note that this is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.
5. Click **Download Log** to download the log too your computer, then you may print or save to a text file.
Click **Clear Event Log** to empty the log.

Configuring event notifications

The server can be configured for automatic notification of selected users when particular types of system events occur. Once warned of an event, these users can run Web Manager to diagnose the problem remotely, with a direct connection or virtual private link to the network.

The event notification can take three forms:

- An *email* message, which the system sends through an SMTP server.
- An *SNMP trap*, to notify a central Network Management Station (NMS) of any events generated by the server; for example, HP OpenView.
- A *syslog* alert enables you to send alerts from a server to a UNIX system log (the UNIX system must have its syslog daemon configured to receive remote syslog messages).



Note: With any of the event notifications, Hitachi Data Systems Support Center recommends setting a notification frequency of Immediately for the most serious alert type (Severe) and to send these alerts to at least two users

Using email alerts

The server can be configured to send emails to specified recipients to alert them on system events. Setting up email alerts requires configuring:

- **SMTP Servers.** The servers on the network to which the reporting server should email alerts.
- **Email Profiles.** Email profiles allow distribution groups to be created, so that email recipients are properly notified based on alert threshold criteria. The server allows classification of email recipients into specific profiles, so that they can receive customized alerts with the depth of focus they require.

For instance, profiles can define different tiers of user responsibility for the server, such that recipients in one profile will only receive alerts on Severe events, while recipients in a second profile receive alerts on Warning and Severe events, and recipients in a third profile get summary emails on all events. In a large user group, dividing these users into separate profiles saves time and simplifies event notification.



Note: A special profile, **SupportProfile**, is intended for support use and sends an alert email to Hitachi Data Systems Support Center. You cannot modify, add recipients to, or remove recipients from the SupportProfile. You can, however, disable this profile. If this profile was changed in an earlier version of the firmware, the only way to reset this profile to its default state is to navigate to the **SMTP Email Profile** page for the SupportProfile. To display this page, go to the **Email Alerts Setup** page (Home > Status & Monitoring > Email Alerts Setup), click details for the SupportProfile to display the **SMTP Email Profile** page for the SupportProfile, then click Restore Default Support Profile.

Procedure

1. Navigate to **Home > Status & Monitoring > Email Alerts Setup** to display the **Email Alerts Setup** page.

Status & Monitoring [Home > Status & Monitoring > Email Alerts Setup](#)

Email Alerts Setup

SMTP Servers


SMTP Primary Server IP/Name: 192.0.2.2
SMTP Secondary Server IP/Name: email.training.bluearc.com
Send Email "From": Group1-Training_TAC_IGNORE


▼ Profile Name	Enabled	Applies To SMU	Immediate Alerts			Summary Alerts			Recipients	
			S	W	I	S	W	I		
<input type="checkbox"/> CustomizedSupportProfile	Yes	No	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	hittrack_email@hds.com	<input type="button" value="details"/>
<input type="checkbox"/> Lab Admin	Yes	No	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	tpillon@bluearc.com	<input type="button" value="details"/>
<input type="checkbox"/> Lab Alerts	Yes	Yes	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	lab.alerts@training.bluearc.com	<input type="button" value="details"/>
<input type="checkbox"/> SupportProfile	No	No	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	hnas.alerts@hds.com	<input type="button" value="details"/>

[Check All](#) | [Clear All](#)

Actions:

Shortcuts: [Configure Email Forwarding](#)

Field/Item	Description
SMTP Primary Server IP/ Name	<p>Enter the host name or IP address of the primary mail server. The server specified as the SMTP server will be used for email alert notification. If the primary SMTP server is offline, the server redirects email notifications to the defined SMTP secondary server.</p> <div>  Tip: As the server should always be in contact with the SMU, Hitachi Data Systems Support Center recommends that the SMU's eth1 IP address be defined as the primary SMTP server. The SMU can be configured for email forwarding and relay any messages to the public mail server. </div>
SMTP Secondary Server IP/ Name	Enter the host name or IP address of the secondary mail server. Email alerts are redirected to this server if the primary SMTP server is unresponsive.
Send Email From	Enter the name you want to appear as the sender (on the From line) of the emails sent according to this profile.
apply	Saves the SMTP server or Send Email From settings specified.
Profile Name	<p>A descriptive name for the profile.</p> <p>A special profile, SupportProfile, is intended for support use and sends an alert email to technical support. You cannot modify, add recipients to, or remove recipients from, the SupportProfile.</p>
Enabled	Indicates whether the profile is enabled.

Field/Item	Description
Applies to SMU	Indicates whether the selected email profile applies to the SMU. (This applies only to an internal SMU.)
Immediate Alerts	<p>A check mark in one of these columns indicates that an immediate email alert is to be sent to the recipients in the named profile when an event of the indicated type occurs.</p> <ul style="list-style-type: none"> • S indicates that an alert will be sent for any severe event. • W indicates that an alert will be sent for any warning event. • I indicates that an alert will be sent for any informational events. <p>It is recommended that immediate email alerts be sent on any severe event.</p>
Summary Alerts	<p>A check mark in one of these columns indicates that a summary email alert is to be sent to the recipients in the named profile when an event of the indicated type occurs.</p> <ul style="list-style-type: none"> • S indicates that an alert will be sent for any severe event. • W indicates that an alert will be sent for any warning event. • I indicates that an alert will be sent for any informational events. <p>It is recommended that immediate email alerts be sent on any severe event.</p>
Recipients	Displays the email addresses that will receive alert emails based on the profile definition.
details	Displays the SMTP Email Profile page, in which you can enable, disable, or edit the email profile.
add	Displays the Add Email Profile page, in which you can create a new email profile.
delete	Deletes the selected email profile.
configure email forwarding	<p>Displays the SMTP Configuration page, which allows you to specify the host name of the email server to which the SMU can send and relay event notification emails.</p> <hr/> <p> Note: This link appears only on clusters with an external SMU</p> <hr/>

2. Specify the SMTP server information in the provided fields.
3. In the **Send Email From** field, specify a name/identifier as the sender. When setting up a new server/cluster, it is important to specify a sender/identifier for the email's From field.
4. Optionally, manage existing email profiles.
 - Click **details** to display the **SMTP Email Profile** page, in which you can enable, disable, or edit the email profile.

- Click **add** to display the **Add Email Profile** page, in which you can create a new email profile.
- Fill the check box for the email profile you want to delete, and click **delete** to remove the selected email profile.
- Click **configure email forwarding** to display the **SMTP Configuration** page, which allows you to specify the host name of the email server to which the SMU can send and relay event notification emails.



Note: This link appears only on clusters with an external SMU.

Daily status emails

A Hitachi NAS Platform system is made up of multiple components. To get an accurate description of the overall status of the various components of the storage system, two daily status emails are generated:

- *Daily status email from the server.* The server's daily status email contains logs of server performance and battery health, descriptive information regarding the health of the server and storage subsystem, and current space utilization by the file systems.
This email is sent to all recipients in all mail profiles in which the `Send a Daily Status Email at midnight` option has been selected.
- *Daily status email from the SMU.* The SMU's daily status email contains a list of the SMU's managed servers and their current firmware versions. It also contains the SMU's current software version. The SMU and server names are links that can be clicked to manage the specified server. The email also provides the ID, model, type (for example, single node or cluster node), and status information about servers.
- *SMU diagnostic emails.* The SMU sends all of its configured email recipients a diagnostic email when any of the following events occur:
 - The server has unexpectedly rebooted.
 - If enabled, once per day at a specified time.

These diagnostic emails contain details regarding the servers, storage, and FC switches managed by the SMU. As the details in these diagnostic mails can be useful to Hitachi Data Systems (should its assistance be required). Hitachi Data Systems also recommends enabling monthly call home emails. When enabled, the SMU sends a full set of server, SMU, and storage diagnostics to Hitachi Data Systems once per month, on a randomly selected day. These provide an archive of the complete configuration of the storage system, which can aid in the detection of problems, provide a wealth of background diagnostic information to Hitachi Data Systems should a problem occur and, if necessary, access to a known good configuration for restoration.



Note: When the monthly diagnostic email is first enabled, an initial email is sent at midnight that night, allowing you to verify that the email configuration is set up correctly.

Adding an email profile

Procedure

1. Navigate to **Home > Status & Monitoring > Email Alert Configuration**, and click **add** to display the **Add Email Profile** page.

[Status & Monitoring](#) [Home](#) > [Status & Monitoring](#) > [Email Alerts Setup](#) > Add Email Profile

Add Email Profile

Settings Used By Server and SMU

Profile Name:

☒ Enable Profile

☒ Send HTML Emails

Add Recipient: **Add**

Recipients: **X**

Server-Specific Settings

☒ Send a Daily Status Email at Midnight ☐ Uuencode Diagnostic Emails Max. Email Length Bytes

☐ Exclude Attachments ☒ Disclose Server Details in Emails

Email Intro Text:

When to Send Alert Emails

Severe: Send Summaries At: hh:mm (24 hour)

Warning: None hh:mm

Information: ☒ Send Empty Summary Alert Emails

☒ Ignore NDMP Events in Immediate Emails

SMU-Specific Settings

☐ Use this profile as the SMU's profile

Custom From Address:

(Default 'from' address is: root@group1-smu.training.bluearc.com)

Send Daily Summary At: hh:mm (24 hour)

To disable 'SMU Daily Status' emails, empty this field.

☐ Enable Call Home Emails (recommended)

Send to: hnas.alerts@hds.com

OK **cancel**

Field/Item	Description
Settings Used By Server and SMU	
Profile Name	Select a name for the profile being created.
Enable Profile	Fill the check box to enable the profile, or leave it inactive.
Send HTML Emails	Fill this check box to receive emails in HTML format. HTML emails are easier to read compared to plain text mails, and this provides easy access to the web UI, because the server name in the email is clickable.
Add Recipient	Enter the email address of the recipient to be added to the profile. Click Add to add the specified recipient to the list of recipients for this email profile.

Field/Item	Description
Recipients	Displays a list of email addresses that will receive emails based on this profile. To delete a recipient from the list, select the email address and click X .
Server-Specific Settings	
Send a Daily Status Email at Midnight	By default, the Send a Daily Status Email at Midnight check box is filled. Detailed emails containing logs of server performance and battery health, descriptive information regarding the health of the server and storage subsystem, and the current space utilization of the file systems will be sent to the specified recipient. To avoid sending daily status emails, clear the check box.
UUencode Diagnostic Emails	Fill this checkbox to uuencode the email attachments sent with the mail that the server automatically sends when it restarts after an unplanned shutdown. This message contains diagnostic information that may help recipients to identify the cause of the problem. By uuencoding the message any virus scanning software at the recipient's site will be bypassed.
Max. Email Length	Limit the size of the email by specifying the maximum number of bytes it can contain. It must be stated numerically, such as: 32768.
Exclude Attachments	Fill this check box to prevent attachments from being sent when daily summary emails are sent.
Disclose Server Details in Emails	By default, the Disclose Server Details in Emails check box is filled. Detailed emails containing restricted or confidential information (account names, IP addresses, portions of user data, and so forth) will be sent to the specified recipient. To avoid sending detailed emails, clear the check box.
Email Intro Text	Custom text to add to the body of the email. You can use this text field to add information or comments to the body of the email. If you are sending HTML emails, you can add basic HTML formatting (italics, bold, new lines and paragraphs, and so on) to the email, and the additional text will be displayed according to the formatting you entered.
When to Send Emails	
Severe/Warning/Information	Select the preferred option for the chosen recipient from the menu: <ul style="list-style-type: none"> • Immediately • Summary • Never
Send Summaries At	Set the time when summary emails should be sent. Set the exact time (<i>hh:mm</i>) in a 24-hour format (for example, 2 PM will be set as 14:00). A second summary can also be sent by entering a time in the second box.

Field/Item	Description
Send Empty Summary Alert Emails	By default, this check box is filled, meaning that empty summary alert emails will be sent to the specified recipient. To avoid sending empty summary emails, clear the check box.
Ignore NDMP Events in Immediate Emails	Fill this check box to prevent emails from being sent when events are generated by the NDMP backup system.
SMU-Specific Settings	
Use this profile as the SMU's profile	Fill the check box to use this profile as the SMU's profile.
Custom From Address	For emails that will be sent by this SMU, enter the address that you want listed as the sender's email address. Note that this field is not available for internal SMUs.
Send Daily Summary At	Set the time when SMU daily status emails should be sent. Set the exact time (<i>hh:mm</i>) in a 24-hour format (i.e. 2 PM will be set as 14:00). To avoid sending daily status emails, empty the field.
Enable Monthly Call Home Emails	Fill the check box to enable monthly call home emails. Clear the check box if you do not to receive monthly call home emails.
OK	Saves the email profile, as specified above.
cancel	Returns to the Email Alerts Setup page without saving the profile.

2. Verify your settings, and click **OK** to save, or **cancel** to decline.

Managing email alerts and profiles

The **Email Alerts Setup** page can be used to delete a profile or modify its properties.

Procedure

1. Navigate to **Home > Status & Monitoring > Email Alerts Setup**, select a profile, and click **details** to display the **SMTP Email Profile** page.

Status & Monitoring [Home](#) > [Status & Monitoring](#) > [Email Alerts Setup](#) > SMTP Email Profile

SMTP Email Profile

Settings Used By Server and SMU

Profile Name: CustomizedSupportProfile

☒ Enable Profile

☒ Send HTML Emails

Add Recipient:

Recipients: hitrack.email@hds.com

Server-Specific Settings

☒ Send a Daily Status Email at Midnight ☐ Uuencode Diagnostic Emails Max. Email Length 32768 Bytes

☐ Exclude Attachments ☒ Disclose Server Details in Emails

Email Intro Text:

When to Send Alert Emails

Severe: Send Summaries At: 08:30 hh:mm (24 hour)

Warning: None hh:mm

Information: ☒ Send Empty Summary Alert Emails

☒ Ignore NDMP Events in Immediate Emails

SMU-Specific Settings

☐ Use this profile (CustomizedSupportProfile) as the SMU's profile

Custom From Address:

(Default 'from' address is: root@group1-smu.training.bluearc.com)

Send Daily Summary At: 08:30 hh:mm (24 hour)

To disable 'SMU Daily Status' emails, empty this field.

☐ Enable Call Home Emails (recommended)

Send to: hnas.alerts@hds.com

Field/Item	Description
Profile Name	Select a name for the profile being created.
Enabled	Fill the check box to enable the profile, or leave it inactive.
Uuencode Diagnostic Emails	Fill this check box to uuencode email attachments. By uuencoding the message, any virus scanning software at the recipient's site will be bypassed.
Send HTML Emails	Fill this check box to receive emails in HTML format. HTML emails are easier to read compared to plain text mails, and this provides easy access to the web UI, because the server name in the email is clickable.
Send Empty Emails	By default, the Send Empty Emails check box is filled. Empty summary emails will be sent to the specified recipient when this is selected. To avoid sending empty summarized emails, clear the check box.

Field/Item	Description
Disclose Email Details to the recipient	By default, the Disclose Email details to the recipient check box is filled. Detailed emails containing restricted or confidential information (account names, IP addresses, portions of user data, and so forth) will be sent to the specified recipient. To avoid sending detailed emails, clear the check box.
Send a Daily Status Email	By default, the Send a Daily Status Email check box is filled. Detailed emails containing logs of server performance and battery health, descriptive information regarding the health of the server and storage subsystem, and the current space utilization of the file systems will be sent to the specified recipient. To avoid sending daily status emails, clear the check box.
Ignore NDMP events in immediate emails	Fill this check box to prevent emails from being sent when events are generated by the NDMP backup system.
Exclude Attachments in Daily Summary Emails	Fill this check box to prevent attachments from being sent when daily summary emails are sent.
Max. Email Length	Limit the size of the email by specifying the maximum number of bytes it can contain. It must be stated numerically, such as: 32768.
When to Send Emails	
Severe/Warning/Information	Select the preferred option for the chosen recipient from the menu: <ul style="list-style-type: none"> • Immediately • Summary • Never
SMU-Specific Settings	
Use this profile as the SMU's profile	Fill the check box to use this profile as the SMU's profile.
Send Email From	For emails that will be sent by this SMU, enter the address that you want listed as the sender's email address. Note that this field is not available for internal SMUs.
Enable Monthly Call Home Emails	Fill the check box to enable monthly call home emails. Clear the check box if you do not to receive monthly call home emails.
Send Summaries At	Set the time when summary emails should be sent. Set the exact time (<i>hh:mm</i>) in a 24-hour format (for example, 2 PM will be set as 14:00). A second summary can also be sent by entering a time in the second box.

Field/Item	Description
Email Intro Text	Custom text to add to the body of the email. You can use this text field to add information or comments to the body of the email. If you are sending HTML emails, you can add basic HTML formatting (italics, bold, new lines and paragraphs, and so on) to the email, and the additional text will be displayed according to the formatting you entered.
Recipients	Displays the current recipient's email address.
Add Recipient	Enter the email address of the recipient about to be added to the profile. Click Add to add the specified recipient to the current profile. Click X to delete the selected recipient from the current profile.

2. Modify the profile by selecting the desired alert options from the menus and check boxes.
3. Verify your settings, and click **OK** to save, or **cancel** to decline.



Note: The SMTP primary server IP/name should always point to SMU's private IP address, while the SMTP secondary server IP/name should always point to the company's main SMTP server.

Using SNMP and syslog

The Simple Network Management Protocol (SNMP) is a standard protocol for managing connected network devices. An SNMP agent can be set up so that Network Management Stations (NMS) or SNMP managers can access its management information.

The server supports SNMP versions 1 and 2c.

SNMP statistics

SNMP statistics (per port and overall in 10-second time slices) are available for activity since the previous reboot, or since the point when statistics were last reset.

Management information base (MIB)

The SNMP agent maintains a Management Information Base (MIB) that is organized in a treelike structure, with each item of data having a unique object identifier (OID) that is written as a series of numbers separated by dots.

The storage server SNMP agent not only supports the MIB-II specification as described in RFC1213, but also provides an Enterprise MIB module, making management facilities available beyond those in the MIB-II specification.

Download the Enterprise MIB module from the SMU (Home > Documentation > SNMP MIB Modules), or contact Hitachi Data Systems Support Center for the latest Enterprise MIB module. The Enterprise MIB module is compiled for SNMP v2c, and is defined in two modules, BLUEARC-SERVER-MIB and BLUEARC-TITAN-MIB.



Note: The NAS server MIB provides information about its own hardware and software. The server MIB cannot provide information about other external hardware, including RAID controllers, physical disks, FC switches, and so on. Those devices provide their own MIBs to monitor such hardware.

Implementing SNMP security

The SNMP agent is provided for monitoring purposes only; it provides read-only access. By default, the SNMP agent does not permit access to the management information base (MIB). Access is enabled by specifying:

- The version of the SNMP protocol with which requests must comply.
- The community names of the SNMP hosts and their associated access levels.
- The IP address or name of hosts from which requests can be accepted (or just choose to accept requests from any host).

Procedure

1. Navigate to **Home > Server Settings > SNMP Access Configuration**.

Server Settings [Home](#) > [Server Settings](#) > SNMP Access Configuration

SNMP Access Configuration

☐ Send traps upon authentication failure
☐ Disable agent
☐ Process SNMPv1 requests only
☐ Process SNMPv2c requests only
☒ Process SNMPv1 and SNMPv2c requests

Accept SNMP Packets On Port:


Send Traps To Port:


☒ Restrict Access To Allowed Hosts

Allowed Hosts: Add Delete

Allowed Communities: Add Delete

apply

Field/Item	Description
Send traps upon authentication failure	Fill this check box if the SNMP agent is to send a trap in the event of an authentication failure (caused, for example, by the SNMP host using an incorrect community string when formulating a request).
SNMP Protocol Support	Using the radio buttons at the top of the page, select the version of the SNMP protocol with which hosts must comply when sending requests to the agent, or alternatively, disable the SNMP agent.
Accept SNMP Packets On Port	Enter the port number that the server monitors for communication through the SNMP protocol.
Send Traps to Port	Enter the port number that the server uses to send traps.
Restrict Access To Allowed Hosts	Fill this check box to restrict protocol access to the hosts specified on this page. Empty the checkbox to enable the protocol to access any host.
Allowed Hosts	<p>To permit requests from authorized hosts only, type the IP address of a host in this field, then click Add to include it in the list. (If the system has been set up to work with a name server, you can type the name of the SNMP manager host rather than its address.)</p> <hr/> <div>  Note: If access is restricted to specified hosts, add the SMU as an allowed host. </div> <hr/>

Field/Item	Description
	To remove a host from the list, select the host you want to remove, then click Delete .
Allowed Communities	<p>Type the name of a community (a password) that will provide authentication into the MIB, and then click Add to include it in the list. Community names are case-sensitive.</p> <hr/> <p> Note: You should define at least one community entry.</p> <hr/> <p>To remove a community from the list, select the host you want to remove, then click Delete.</p>
apply	Click apply to save configuration changes and close the page.

2. Verify your settings, and click **apply** to save your changes.

Sending SNMP traps

A trap is unsolicited information that the SNMP agent sends to a manager. It enables the agent to alert the manager to some unusual system event. The SNMP agent supports the following limited set of traps:

- **AuthenticationFailure.** Indicates that the SNMP agent received a request from an unauthorized manager. Either the manager used an incorrect community name or the agent has been set up to deny access to the manager.
- **ColdStart.** Indicates that the SNMP agent has started or been restarted.
- **LinkUp.** Indicates that the status of an Ethernet link has changed from *Down* to *Up*.

Procedure

1. Navigate to **Home > Status & Monitoring > SNMP Traps Setup**.

Field/Item	Description
Notification Frequency	<p>Using the list, select the notification frequency for each type of alert:</p> <ul style="list-style-type: none"> • Severe Alerts: The specified component has failed in a way that poses a significant threat to the continued operation of the system. • Warning Alerts: The specified component needs attention but does not necessarily represent an immediate threat to the continued operation of the system. • Information Alerts: The specified component is operating normally and is not displaying an alarm condition.
Trap Recipients	<p>In this area, enter the hosts to which this server will send traps. In the Host field, enter the IP address of an SNMP host to associate with each community. (If the system has been set up to work with a name server, you can type the name of the SNMP manager host rather than its address.)</p> <p>In the Community field, type the name of the SNMP community (community names are case-sensitive).</p> <p>Click Add to save the information in the list.</p> <p>You can delete an entry in the list by selecting it and clicking the X.</p>
apply	Saves the settings.

2. Enter the necessary information, and click **apply**.

Configuring syslog notifications

You can use syslog notification to send a syslog alert from the server to a UNIX system log when three types of events occur. The UNIX system must have its syslog daemon configured to receive remote syslog messages.

Procedure

1. Navigate to **Home > Status & Monitoring > Syslog Alerts Setup**.

Syslog Alerts Setup

Notification Frequency

Severe Alerts:

Warning Alerts:

Information Alerts:

Syslog Servers

Field/Item	Description
Notification Frequency	Select the notification frequency for each type of alert: <ul style="list-style-type: none">• Severe Alerts: The specified component has failed in a way that poses a significant threat to the continued operation of the system.• Warning Alerts: The specified component needs attention but does not necessarily represent an immediate threat to the continued operation of the system.• Information Alerts: The specified component is operating normally and is not displaying an alarm condition.
New Syslog Recipient	In this area, enter the syslog servers to which this server will send alerts. In the first field, enter the IP address or host name of the syslog server, and click Add to save the address in the list. You can delete an entry in the list by selecting it, and clicking the X .
apply	Click to apply and save the configuration.

2. Enter the necessary information.

Testing alert configurations

After setting up the alert configuration, send a test alert to all selected recipients.

Procedure

1. Navigate to **Home > Status & Monitoring > Send Test Event**.
2. Select the type of message to send from the list, enter a test message, and click **test**.

Clearing logs with Windows Event Viewer

The Windows Event Viewer utility allows clearing (deleting) events related to specific categories, such as 'system' or 'security.'

On a Windows server, when an administrator clears a particular category of events, only events in that category are deleted.

For the NAS server, however, when managing the NAS server event log using the Windows Event Viewer, you cannot delete only a particular category of events. If you try to clear only security or system event entries from the log, you actually clear all events (security and system) from the NAS server's event log.

File system auditing

File system auditing monitors and records file access and deletion operations performed through the CIFS protocol. These operations are recorded in the file system's audit log. You can then display the file system's audit log, and use a remote Windows Event Viewer to save the log entries for later review. File system audit logging is performed and controlled on a per file system basis.

Because CIFS defines open and close operations, auditing file system object access performed by clients using other protocols would be costly in terms of system performance, because each I/O operation would have to be audited as an open operation. Therefore, when file system auditing is enabled, by default, only clients connecting through the CIFS protocol are allowed access to the file system. Access by clients using other protocols, like NFS, can, however, be allowed. When such access is allowed, access to file system objects through these protocols is not audited.



Note: You can configure file system auditing to deny access to clients connecting with protocols that cannot be audited (NFS).

After a file has been externally migrated (migrated to an external server), for example to a Hitachi Content Platform (HCP) system, subsequent access to the file through the NAS server is audited as if the file were still local.

For known users (users with a Windows user mapping), the NAS server logs Object Access events 560, 562, 563 and 564. As with the Windows operating system, auditable events for objects are specified by SACLs (system access control lists). Auditing events are logged under the following conditions:

- 560 – open handle
This event is logged when a network client asks for access to an object. An access check is performed against the DACL (discretionary access control list) and an audit check is performed against the SACL. If the result of the access check matches the result of the audit check, an audit record is generated.
- 562 – close handle
This event is logged when an application closes (disposes of) an existing handle, and is logged in conjunction with event 560.
- 563 – open handle for delete
This event is logged when a network client asks for access to a file using the CreateFile call, and the delete-on-close flag is specified. An access check is performed against the DACL and an audit check is performed against the SACL. If the result of the access check matches the result of the audit check, an audit record is generated.
For successful deletions, the audit records the accesses that were granted, and for failures the audit records the accesses that were requested.
- 564 – delete
This event is logged when an application closes (disposes of) an existing handle, and is logged in conjunction with event 563.

About file system audit logs

The file system audit log is buffered in memory, and may be permanently stored in a file in the file system being audited. Active audit log files are stored in Windows event log file format (.evt) so that standard tools can access them. The name, location, size of the active audit log file, log file retention, and active log file backup settings are defined when enabling auditing for a file system.



Note: File System Audit logs are saved in Windows XP format. An effect of this is that, depending upon how the saved .evt file is opened, a Windows Vista or Windows 2008 Server event viewer can report the file as corrupted, or might not be able to fully interpret the events. Note that the same situation occurs when a Windows Vista event viewer is used to display saved logs from an XP system. To display the logs correctly, use a Windows XP event viewer.

Audit log files are limited in size, and the retention behavior when a log fills is configurable. When an audit log reaches its maximum size, log entries (file system events) can be overwritten, or the full audit log can be saved, and a new log started



Note: All file system audit log parameters are specified on a per file system basis.

You can specify a backup policy, which backs up the active log at regular intervals, and starts a new active log file. Backup log files are created in the same directory as the active audit log file.

In the event of a server crash, active file system audit logs are recovered only if a rollback is performed on restart. Note that a rollback may reset the audit log file to a time when it can be recovered, thus saving some records that would otherwise be lost.

Controlling file system auditing

File system auditing requires that a file system audit policy be defined for the file system to be monitored, and that auditing is enabled for the specific file system. File system auditing is performed and controlled on a per file system basis.

Creating a file system audit policy

The file system audit policy specifies access restrictions for clients connecting through unauditable protocols (if access is allowed or denied), and specifies audit log details. The audit log policy specifies naming, location in the file system, size, the log roll over policy, and the backup policy.

Procedure

1. Navigate to **Home > Files Services > File System Audit Policies**, and click **add** to display the **Add File System Audit Policy** page.

File Services [Home](#) > [File Services](#) > [File System Audit Policies](#) > Add File System Audit Policy

Add File System Audit Policy

EVS / File System: donotdelete / donotdelete [change...](#)

Access via Unsupported Protocols

☒ **Deny Access**
Client access to the file system via un-auditable protocols (such as NFS) will be denied; refer to Help for more information

☐ **Allow Access (without auditing)**
Allow access but do not create any auditing events for un-auditable protocols (such as NFS)

Audit Log

Active Log File Name:
(File name entered must have .evt extension)

Logging Directory: [browse...](#)
(Directory will be created if it does not exist)

Maximum Log File Size:

Log roll over policy

☒ New
☐ Wrap

Backup Policy

Backup Interval: minutes

Number of files to retain:

[OK](#) [cancel](#)

Field/Item	Description
EVS/File System	Lists the currently selected EVS and file system, to which the audit policy will apply. Click change to go to the Select a File System page, where you can select a different EVS and file system.
Access via Unsupported Protocols	When clients attempt to access the file system through a protocol that does not support auditing (such as NFS), this setting determines if those clients are permitted to access the file system. You can select either: <ul style="list-style-type: none"> • Deny Access. Client access to the file system using un-auditable protocols (such as NFS) is denied. • Allow Access. Allows client access to the file system using un-auditable protocols (such as NFS), but does not create any auditing events.
Active Log File Name	Specify the file name for the file system audit log. The file name must have an .evt extension. The default file name is audit.evt.
Logging Directory	Specify the directory within the file system in which the file system audit log files are saved. You can use the browse

Field/Item	Description
	button to search for an existing directory, or enter the name of a directory to be created.
Maximum Log File Size	Specify the maximum size of the active audit log file in KB, MB. The maximum log file size is 50 MB.
Log roll over policy	Determines what the system does once the active audit log file is full (when it reaches the Maximum Log File Size). You can select either: <ul style="list-style-type: none"> • Wrap, which causes the system to delete the oldest existing audit entry to allow room for a new entry. • New, which causes the system to create a new active audit log file. The default is New.
Backup Interval	Specify the time (in minutes) between automatic backups of the active audit log. The backup interval must be between 5 and 14400 minutes (10 days). A value of 0 disables the automatic backups. The default is 0.
Number of files to retain	Specify the number of backup audit log files to retain. The default is 10 .
OK	Saves the file system audit policy.
cancel	Exits this page without creating the file system audit policy.

- Specify the access settings for unsupported (unauditable) protocols
 - **Deny Access.** Client access to the file system using unauditable protocols (such as NFS) is denied.
Specifying **Deny Access** generates an error if there is an NFS export for the file system or the file system has a FTP user that has a directory available. To ensure this error is not generated, you can remove the NFS export for the file system, remove the FTP user, or select the **Allow Access** option.
 - **Allow Access.** Allows client access to the file system using unauditable protocols (such as NFS), but does not create any auditing events.
- Specify the name for the active audit log file. The file type suffix must be `.evt`.
- Click **browse** to specify an existing logging directory, or enter the name of a directory to create.
For ease of access to the audit log files, the logging directory should be within in a CIFS share that can be accessed by those who need to review the access log.
- Specify the maximum log file size.
- Specify the roll over (retention) policy.
- Specify the backup interval.
- Specify the number of files to retain.
- Click **OK** to save the policy as specified.

Enabling auditing for a file system

File system auditing can be enabled on a per-file system basis.

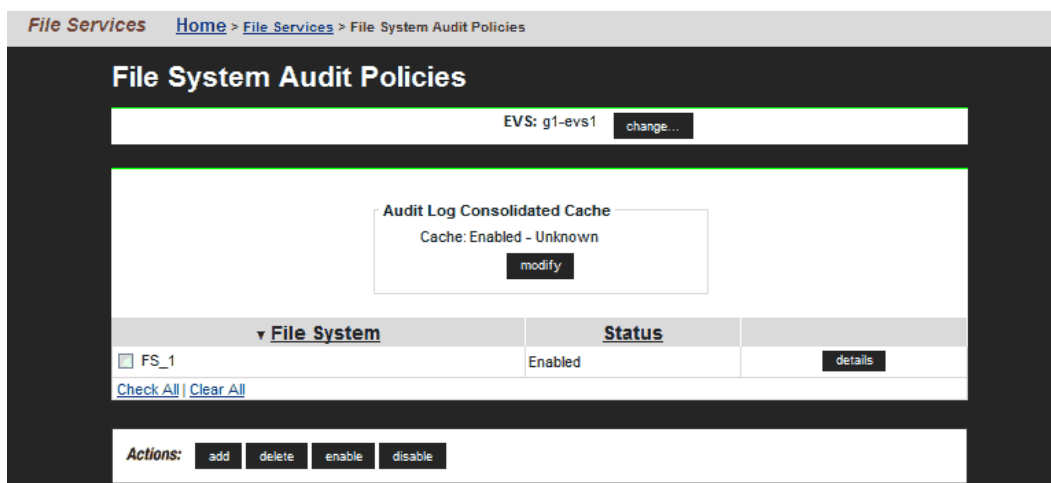


Note: By default, when file system auditing is enabled, access to the file system will be limited to the CIFS protocol. Access by clients using other protocols, like NFS, can, however, be allowed. When such access is allowed, access to file system objects through these protocols is not audited.

To enable file system auditing for a particular file system, the file system must be added to the file system audit list.

Procedure

1. Navigate to **Home > File Services > File System Audit Policies**.



Field/Item	Description
EVS	Lists the EVS to which host the file system is assigned. Click change to go to the Select an EVS page, where you can select a different EVS.
File System	Lists all file systems in the specified EVS that have an audit policy.
Audit Policy Status	Indicates whether file system auditing is enabled or disabled.
details	Displays the File System Audit Policy Details page, in which you can change the auditing options for a file system.
add	Displays the Add File System Audit Policy page, in which you can set the auditing options for a file system. Only one audit policy is allowed per file system.
delete	Deletes the audit policy for a selected file system.
enable	Enables file system auditing for the selected file system.
disable	Disables file system auditing for the selected file system.

2. If the file system on which you want to enable auditing is listed, an audit policy has already been defined for that file system.
 - If the Audit Policy Status is enabled, logging is already enabled for the file system, and no further actions are required.
 - If the Audit Policy Status is disabled, fill the check box next to the file system name, and click **enable**.

If the file system on which you want to enable auditing is not displayed, a file system audit policy may not have been defined for that file system, or the file system may have an audit policy defined, but the file system is not in the currently selected EVS.

3. Click **change** to display the **Select an EVS** page, in which you can select a different EVS.
 - If, after selecting the EVS that hosts the file system, the file system on which you want to enable auditing is now listed on the **File System Audit Policies** page, fill the check box next to the file system name, and click **enable**.
 - If, after selecting the EVS that hosts the file system, the file system on which you want to enable auditing is still not displayed, you must define a file system audit policy for that file system. Click **add** to display the **Add File System Audit Policy** page, in which you can set the auditing options for a file system.

Modifying a file system audit policy

Procedure

1. Navigate to **Home > Files Services > File System Audit Policies**.

If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, in which you can select a different EVS.

- Click the **details** button on the file system with the audit policy you want to modify to display the **File System Audit Policy Details** page.

File Services [Home](#) > [File Services](#) > [File System Audit Policies](#) > File System Audit Policy Details

File System Audit Policy Details

File System: FS_1
Auditing: Enabled disable

Access via Unsupported Protocols

☒ Deny Access
Client access to the file system via un-auditable protocols (such as NFS) will be denied; refer to Help for more information

☐ Allow Access (without auditing)
Allow access but do not create any auditing events for un-auditable protocols (such as NFS)

Audit Log

Active Log File Name:
(File name entered must have .evt extension)

Logging Directory: browse...
(Directory will be created if it does not exist)

Maximum Log File Size: KB

Log roll over policy

☒ New
☐ Wrap

Backup Policy

Backup Interval: minutes

Number of files to retain:

OK cancel

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	Lists the currently selected EVS and file system, to which the audit policy will apply. Click change to go to the Select a File System page, where you can select a different EVS and file system.
Auditing	Indicates whether file system auditing is enabled or disabled. Click enable or disable to toggle the auditing mode.
Access via Unsupported Protocols	When clients attempt to access the file system through a protocol that does not support auditing (such as NFS), this setting determines if those clients are permitted to access the file system. You can select either: <ul style="list-style-type: none"> Deny Access. Client access to the file system using un-auditable protocols (such as NFS) is denied. Allow Access. Allows client access to the file system using un-auditable protocols (such as NFS), but does not create any auditing events.
Active Log File Name	Specify the file name for the file system audit log. The file name must have an .evt extension. The default file name is audit.evt.

Field/Item	Description
Logging Directory	Specify the directory within the file system in which the file system audit log files are saved. You can use the browse button to search for an existing directory, or enter the name of a directory to be created.
Maximum Log File Size	Specify the maximum size of the active audit log file in KB, MB, or GB. The default is 512KB.
Log roll over policy	Determines what the system does once the active audit log file is full (when it reaches the Maximum Log File Size). You can select either: <ul style="list-style-type: none"> • Wrap, which causes the system to delete the oldest existing audit entry to allow room for a new entry. • New, which causes the system to create a new active audit log file. The default is New.
Backup Interval	Specify the time (in minutes) between automatic backups of the active audit log. The backup interval must be between 5 and 14400 minutes (10 days). A value of 0 disables the automatic backups. The default is 0.
Number of files to retain	Specify the number of backup audit log files to retain. The default is 10 .
OK	Saves the file system audit policy.
cancel	Exits this page without creating the file system audit policy.

3. Modify the policy as required.
4. Click **OK** to save the policy as specified.

Enabling or disabling auditing for a file system

Procedure

1. Navigate to **Home > Files Services > File System Audit Policies**.
If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, in which you can select a different EVS.
2. Fill the check box next to the name of the file system with the audit policy you want to enable or disable.
3. Click **Enable** to allow a disabled policy to function again, or click **Disable** to stop the policy from functioning.
When disabled, file system access operations are not logged, and protocol restrictions are not enforced. Note that disabling a policy does not delete it.

Deleting a file system audit policy

Procedure

1. Navigate to **Home > Files Services > File System Audit Policies**.
If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, in which you can select a different EVS.
2. Fill the check box next to the name of the file system with the audit policy you want to delete, and click **delete**.



Note: Existing log files are not deleted automatically when a policy is deleted. If you want to delete these logs, you must do so manually,

Displaying file system audit logs

The NAS server supports using a remote Windows Event Viewer to display file system audit log events. The audit log files are shown in the "FS" (file system) log, which can be displayed by the Windows Event Viewer, assuming that:

1. You have used the `audit-log-consolidated-cache` command to configure a single consolidated cache file (the `audit-log-consolidated-cache`).
If the cache file is not configured, the Windows Event Viewer cannot view file system events. The consolidated cache file has a default size of 10MB, and a maximum size of 50MB.



Note: Only one consolidated cache file can be configured per EVS. Audit events from all file systems assigned to that EVS are collected into this single consolidated cache file.

When you create the consolidated cache file, you must specify the name of the file system in which the file will be stored. The cache file is located in the `.audit` directory of the root of the named file system. The default name for the consolidated cache file is `audit_cache.evt` (audit log files for individual file systems have a default name of `audit.evt`).

2. The logging directory is within a CIFS share.

Using the Windows Event Viewer, you can display, save, and clear the local event logs, or those on a remote computer. Audit logs can be saved in several formats, including a `.evt` event format or a plain text file. The Windows Event Viewer can only save in `.evt` format to a file on the same computer as the event log, because it is the computer being viewed that does the copy

(meaning the Event Viewer does not just read the event log and write it to a file). The Event Viewer can also be used to open and display saved audit log files.

Optionally, you can send file system audit logs to a remote syslog server using the `audit-syslog` command. Enter `man audit-syslog` at the CLI, or see the *Command Line Reference* for more information.

FTP auditing

FTP audit logging is controlled on a per-EVS basis. When enabled, the system maintains an audit log which tracks user activity performed through the FTP protocol for all file systems in the EVS. Each time a user takes any of the following actions, the system records the event:

- Logging in or out (including when a session timeout occurs).
- Renaming or deleting a file.
- Retrieving, appending or storing a file.
In this case, the system records the success or otherwise of the action at both its start and end.
- Creating or removing a directory.

Displaying FTP Audit Logs page

The **FTP Audit Logs** page displays the FTP audit logging status for each EVS in the server or cluster. Using this page, you can view FTP logging status, enable or disable FTP audit logging, and you can also display the **FTP Audit Log Details** page, which allows you to configure log file details.

Procedure

1. Navigate to **Home > File Services > FTP Audit Logs**.

File Services [Home](#) > [File Services](#) > FTP Audit Logs

FTP Audit Logs

Show 20 items per page

▼ EVS	File System	Path	Status	
<input type="checkbox"/> g1-ews3	*Unknown*		Disabled	details
<input type="checkbox"/> donotdelete	*Unknown*		Disabled	details
<input type="checkbox"/> g1-ews1	*Unknown*		Disabled	details
<input type="checkbox"/> LNAS	*Unknown*		Disabled	details
<input type="checkbox"/> g1-ews2	*Unknown*		Disabled	details
<input type="checkbox"/> EVS1	*Unknown*		Disabled	details

[Check All](#) | [Clear All](#)

Actions: [enable](#) [disable](#)

For each EVS in the server or cluster, this page lists the status of FTP audit logging, and displays the file systems being monitored, as well as the path to the FTP audit logs for each monitored file system.

Field/Item	Description
EVS	Lists the file serving EVS in the server or cluster.
File System	Lists the file systems in the server or cluster.
Path	Displays the directory path in the file system where the FTP audit log is located.
Status	Indicates whether FTP auditing is enabled or disabled.
details	Display the FTP Audit Log page, which allows you to configure FTP audit logging for the file system.
enable	Fill the check box for an EVS, and click enable to enable FTP auditing for the EVS.
disable	Fill the check box for an EVS, and click disable to disable FTP auditing for the EVS.

Enabling or disabling FTP audit logging for an EVS

FTP Audit Logging is enabled or disabled on a per-EVS basis, meaning that is enabled or disabled for all file systems served by the EVS and accessed through the FTP protocol.

Procedure

1. Navigate to **Home > File Services > FTP Audit Logs**
2. Fill the check box for the EVS for which you want to enable or disable FTP audit logging.

- If FTP Audit Logging is disabled, you can enable it by clicking **enable**.
- If FTP Audit Logging is enabled, you can disable it by clicking **disable**.

Configuring FTP audit logging

Procedure

1. Navigate to **Home > File Services > FTP Audit Logs**.

- Click **details** to display the **FTP Audit Log Details** page for the EVS for which you want to configure FTP audit logging.

[File Services](#) [Home](#) > [File Services](#) > [FTP Audit Logs](#) > FTP Audit Log Details

FTP Audit Log Details

EVS: g1-evs3

Audit Logging: Disabled

File System: No File System Selected

Logging Directory:



Path Options

This option applies only when 'File System' or 'Logging Directory' values are changed.

☒ Create path if it does not exist. (See online help for security implications).

Max. Records per Log File:

Max. Number of Log Files:

Field/Item	Description
EVS	Lists the currently selected EVS and file system, to which the audit configuration will apply.
Audit Logging	Indicates if FTP audit logging is enabled or disabled for the EVS.
File System	Displays the name of the file system that will contain the FTP audit log files. Click change to select a different file system.
Logging Directory	<p>Displays the directory path in the file system in which the FTP audit log files are stored. The path options allow you to select an existing directory, or to create the directory if it does not already exist.</p> <hr/> <p> Note: The browse... button only exists if the path being created is the path in a file system, not a namespace.</p> <hr/> <p> Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users; for example: the permissions are set to rwxxrwxrwx. It is recommended that such directories are created using CIFS or NFS, or that such directories are given the desired permissions explicitly after being created using this option.</p> <hr/>
Max. Records per Log File	Specifies the maximum number of records per log file. Once the maximum number of records per file is reached, a new log file is started. Each log file is a tab-delimited text file containing one line per FTP event. Besides logging the date and time at which an event occurs, the system logs the user name and IP address of the client and a description of the executed command.
Max. Number of Log Files	Specifies the maximum number of log files to be kept. Once the maximum number of log files is reached, when the current log file

Field/Item	Description
	becomes full, the oldest log file is deleted. The newest log file is called <code>ftp.log</code> , and the older files are called <code>ftp<i>n</i>.log</code> (the larger the value of <i>n</i> , the older the file).
OK	Save the configuration.
cancel	Exits without saving the configuration.

3. In the File System field, choose a file system in which to keep the log files. Click **change** to see a list of file systems in the EVS.
For optimum performance, keep the log files on a different system drive than the files that users will access over FTP
4. Specify the logging directory.
The logging directory specifies the location in which the FTP audit logs are kept. In the Logging Directory field, specify the directory in which to keep the log files. Click **browse** to choose an existing directory, or specify a path to be created. To create the path automatically when it does not already exist, fill the check box **Create path** if it does not exist.
5. In the **Max. Number of Records per Log File** field, specify the maximum number of records to store in each log file.
For optimum performance, produce a small number of large files instead of a large number of small files.
6. In the **Max. Number of Log Files** field, specify the maximum number of log files to keep.
Once it has reached this limit, the server deletes the oldest log file each time it creates a new one.
7. Click **OK** to save the FTP audit logging configuration.

Displaying FTP audit logs

FTP audit logs can be displayed with a text editor. If the logging directory is within an NFS export or a CIFS share, access the directory, and open the log file. If the logging directory is available through FTP, you can download the file, then open it with a text editor.

Monitoring Fibre Channel switches

The server allows you to add Fibre Channel (FC) switches to the System Monitor, so you can easily check FC switch connectivity status, which indicates whether the SMU received a response to an Ethernet ping of its last-known IP address. The connectivity status does not indicate whether the FC switch has connectivity with the storage subsystem.

When adding an FC switch to the System Monitor, you can associate it with one or more servers. After an FC switch has been associated with a server, you can monitor switch connectivity status, display log events and SNMP

traps, download FC switch diagnostic information, and configure emailing of switch-related diagnostic information.

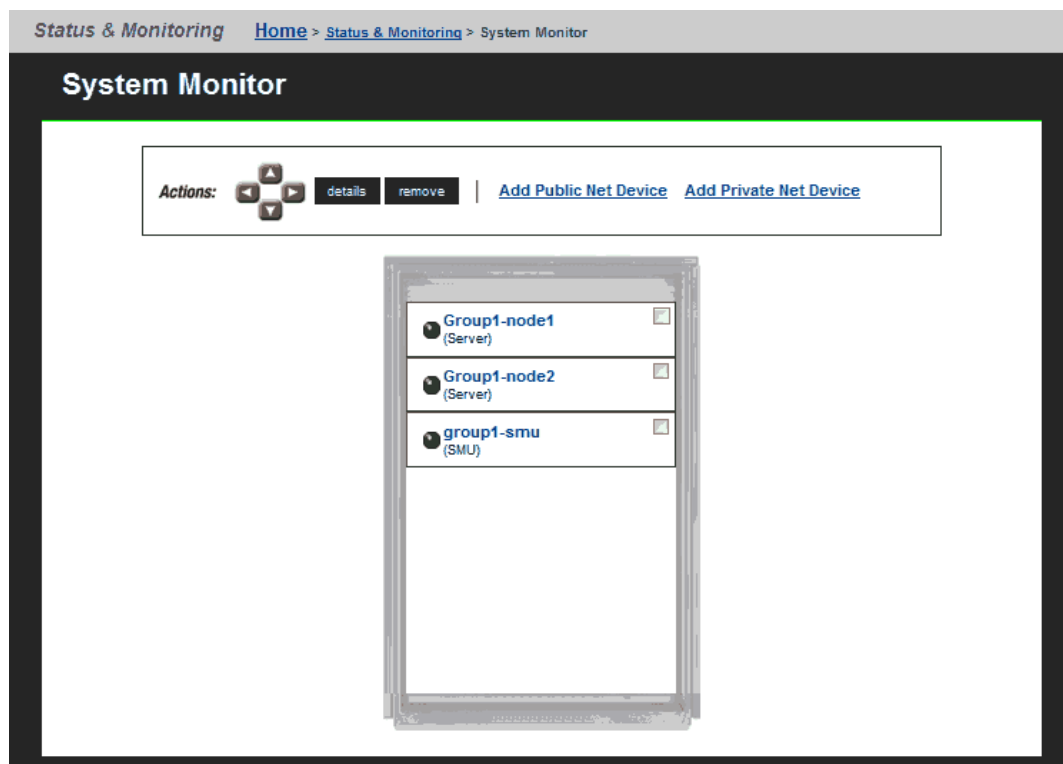
Displaying Fibre Channel switch connectivity status

The System Monitor displays FC switch connectivity status at a glance, and also lists FC switches, which can be selected to display detailed switch information.

Using System Monitor to display switch connectivity status

Procedure

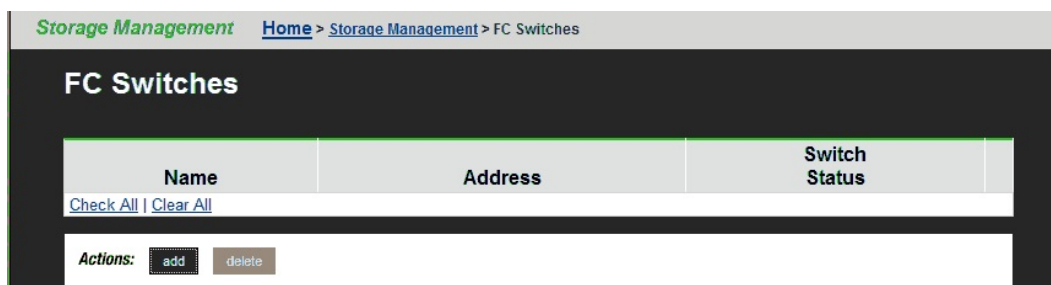
1. Navigate to **Home > Status & Monitoring > System Monitor**.
The status indicator next to the FC switch indicates its connectivity status.



Using Web Manager to display switch connectivity status

Procedure

1. Navigate to **Home > Storage Management > FC Switches**.



Field/Item	Description
Name	The name of the switch, defined when the switch was added. This name should be sufficiently descriptive as to be able to identify the switch.
Address	The IP address or DNS name of the switch, defined when the switch was added.
Switch Status	An indicator of the connectivity status of the switch. Connectivity status indicators are: <ul style="list-style-type: none">• Green – OK. A response was received from a ping of the last-known IP address of the switch.• Gray – Determining state. A FC switch will appear as gray for up to 60 seconds, immediately after it is added. After a ping of the switch IP address, the status will change to OK or severe (green or red), depending on whether there was a response to the ping.• Red – Severe. No response was received from a ping of the IP address of the switch.
details	Displays the FC Switch Details page for the switch. From the FC Switch Details page, you can open the embedded management interface for the switch (if available), and change the switch name or address.
add	Opens the Add FC Switch page.
delete	Deletes one or more selected FC switches.

Adding FC switches


After adding an FC switch, the SMU displays it in the System Monitor, with connectivity status. Because multiple servers or clusters might use the storage connected to an FC switch, it can be associated with multiple servers or clusters managed by an SMU, thereby appearing in the System Monitor for all servers and cluster to which it has been associated.

Procedure

1. Navigate to **Home > Storage Management > FC Switches**, and click **add** to display the **Add FC Switch** page.

2. Enter the requested information.

Field/Item	Description
Associate Existing Switch with name (currently managed server)	Select an existing switch to associate with the named server or cluster. When you associate a switch with a managed server or a cluster, the switch is added to the system monitor of that server/cluster.
Monitor Switch	Use the list to select the switch you want to associate with the named server/cluster.
Add New Switch	Select to add a new FC switch. After the switch has been added, you can associate it with a managed server or a cluster.
Name	The name you want to use to refer to the switch. This name should be sufficiently descriptive as to be able to identify the switch.
Name/IP Address	A fibre channel switch can be specified by IPv4 or IPv6 address, or by a host name. If an IPv6 address is specified, the SMU will only be able to monitor the switch if the SMU is configured with an IPv6 address. Additionally, if the switch is given by host name, and that host name resolves to an IPv6 address, monitoring will only be possible if an IPv6 DNS server is provided.
Username	Enter the user login name for the embedded management interface of the FC switch.
Password	Enter the password associated with the user name for the embedded management interface of the FC switch.
Use http/https/Telnet/other on port...	From the list, select the protocol and port for connecting with the embedded management interface of the FC switch. Defaults are <i>http</i> protocol and <i>port 80</i> .

Field/Item	Description
	 Note: If <i>http</i> , <i>https</i> , or <i>Telnet</i> , clicking the switch in the System Monitor displays the embedded management interface. If <i>other</i> , the FC Switch Details page is displayed instead of the management interface
OK	Saves configuration changes, and closes the page.
cancel	Closes the page without saving configuration.

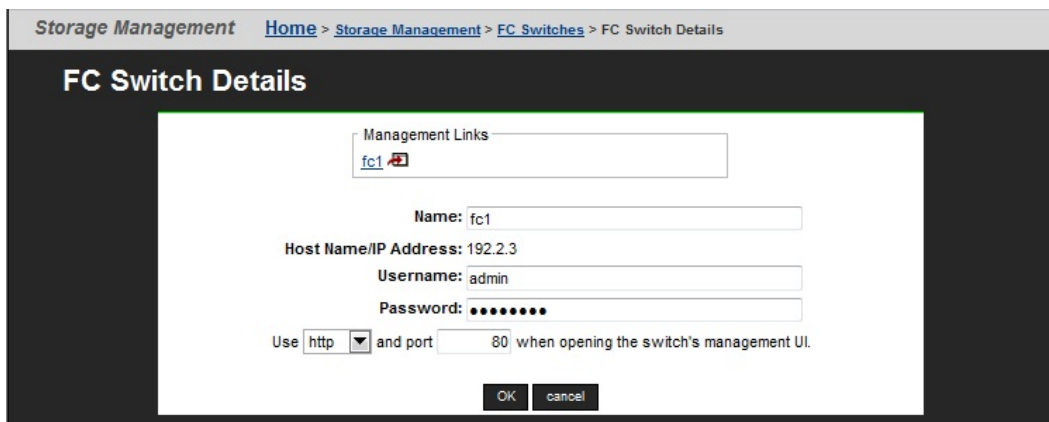
3. Verify your settings, and click **OK** to save, or **cancel** to decline.

Displaying or changing details for an FC switch


You can display a list of the FC switches that have been added to the System Monitor of any server or cluster managed by an SMU on the **FC Switches** page. After you have displayed this list, you can display and change details for a switch.


Procedure

1. Navigate to **Home > Storage Management > FC Switches**, and click detail for a selected switch to display the **FC Switch Details** page, which lists all FC switches that have been added to the System Monitor of the server/cluster.



2. As needed, display or modify the switch information.

Field/Item	Description
Management Links	<p>This area provides links to the embedded management interfaces for the FC switch. Click a link to open the interface.</p> <hr/>  Note: The FC switch management interface might or might not support multiple concurrent logins. Refer to

Field/Item	Description
	the documentation for the switch regarding use of the embedded management interface.
Name	Name of the switch, specified when the switch was added. This name should be sufficiently descriptive as to be able to identify the switch.
Name/IP Address	The IP address or DNS name of the switch, specified when the switch was added.
Username	User login name for the embedded management interface of the FC switch.
Password	Password associated with the user name for the embedded management interface of the FC switch.
Use http/https/Telnet/ other on port...	Protocol and port for connecting with the embedded management interface of the FC switch. Defaults are <i>http</i> protocol and <i>port 80</i> . <div>  Note: If <i>http</i>, <i>https</i>, or <i>Telnet</i>, clicking the switch in the System Monitor displays the embedded management interface. If <i>other</i>, the FC Switch Details page is displayed instead of the management interface. </div>
OK	Saves configuration changes, and closes the page.
cancel	Closes the page without saving configuration changes.

3. Verify your settings, and click **OK** to save, or **cancel** to decline.

Optimizing performance with Performance Accelerator

The Performance Accelerator feature optimizes throughput and IOPS capacity in the NAS Platform system by enabling very-large-scale integration (VLSI) features in the NAS server. Both throughput and IOPS capacities are significantly increased. To maximize throughput in the VLSI, the PCIe connection between the SI fpga and the Tachyon Fibre Channel controller is increased from four to eight lanes. This lane increase doubles the available bandwidth of the connection, providing greater throughput and speed. Performance Accelerator enhances the IOPS component by increasing the number of cache controllers from one to two, within the SI FPGA, maximizing the available amount of cache controller processing power. If a bottleneck previously existed in the PCIe connection to the Tachyon Fibre Channel controller, or to the SI cache controller, Performance Accelerator might reduce or eliminate such a bottleneck.



Note: Performance Accelerator is available only on the NAS Platform 3090 G1 and NAS Platform 3090 G2 servers. Installing Performance Accelerator on other servers has no effect.

Determining if Performance Accelerator will increase system performance

To evaluate the current throughput component, measure the current system throughput. If the current system throughput is close to the "standard" throughput limits, then it is likely that the PCIe connection to the Tachyon Fibre Channel controller is not optimized for performance. Performance Accelerator might bring a performance improvement. The standard read speed, on newer systems equipped with QE4+ Tachyon controllers, is 880 MB/sec; the standard write speed is 800 MB/sec. On older systems, equipped with QX4 Tachyon controllers, the standard read speed is 880 MB/sec; the standard write speed is 640 MB/sec.

For the IOPS component, collect a PIR while the system is under maximum load. Examine the SI utilization by looking at the `"si_busy_clocks_last_second_percentage statistic"` in the `logged-statistics.csv` file. If this file shows that the SI FPGA is very busy (at 90 to 100 percent active, with the standard being 72,000 ops/sec), then it is likely that the SI cache controller is not optimized, and Performance Accelerator might significantly improve performance.

Installing Performance Accelerator

Performance Accelerator is enabled by installing its license.

Testing the Performance Accelerator installation

Performance Accelerator enables additional PCIe lanes in the VSLI to connect to the Tachyon Fiber Channel controller. If these lanes have not been previously tested, the server will perform a full power on self test (POST) to ensure that the lanes are working. If the POST test passes, then both components of Performance Accelerator are enabled when the server boots. If the POST test fails, then only the IOPS (dual cache controller) component of Performance Accelerator is enabled, and an error event is generated.



Note: A full POST test is only possible if there is no stale data in NVRAM left over from deleted file systems that had associated NVRAM content. Stale data is cleared from NVRAM by unmounting file systems thoroughly, and using the `nvpages list` command to inspect for stale data.

Uninstalling Performance Accelerator

Procedure

1. Removing the Performance Accelerator license.
2. Rebooting the server, using the supervisor-level `reboot-app` command. In a cluster, reboot one node at a time.

Troubleshooting Performance Accelerator

At boot time, Performance Accelerator writes the following line to the server dblog:

```
Performance Accelerator: licensed 1, tptelc 1, mtds_passed 0,  
tpcurrent 0, tpprevious 0, dcmode unset
```

The following table defines the meaning of each field in the line:

Field	Description
licensed	1 if Performance Accelerator is licensed, 0 otherwise.
tptelc	1 as long as the throughput component of Performance Accelerator is not disabled by the fci4 telc (see below), 0 otherwise.
mtds_passed	1 if full POST has run and passed, 0 otherwise.
tpcurrent	1 if licensed=1 and tptelc=1 and mtds_passed=1, 0 otherwise.
tpprevious	The value of tpcurrent on the previous boot.
dcmode	The value of the telc used to force dual code mode behavior. If "unset", the default behavior ("striped") is used, as long as Performance Accelerator is licensed.

Verifying that the throughput component of Performance Accelerator is enabled

Use the dev-level **fci-info pciex_status**; for example:

```
mercuryc4(MMB):$ fci-info pciex_status  
fc  
fc                pciex_status = 0x27f00006 670040070  
fc                pciex_num_active_lanes : 0x8 8
```

Note that the output in the last column of the last line, "pciex_num_active_lanes", is "0x8 8", indicating that the PCIe connection between the SI FPGA and the Tachyon Fibre Channel controller is successfully increased from four to eight lanes. The output is "4" if Performance Accelerator is disabled.

Verifying that the IOPS component of Performance Accelerator is enabled

The IOPS component can be verified using the dev-level **si-chip config** command; for example:

```
mercuryc4(MMB):$ si-chip config  
config = 0x29980 170368  
dual_cache_mode : 0x3 3
```

The "dual_cache_mode" shows "3" if Performance Accelerator is enabled ("0" if disabled).

Disabling the throughput component

Procedure

1. Use the dev-level `telcset fci4 true` command.
2. Reboot the server, using the supervisor-level `reboot-app` command. In a cluster, reboot one node at a time.
3. Set the `telc` on all nodes.

Postrequisites

To reenable the throughput component, delete the `fci4 telc`, and reboot.

Disabling the IOPS component

Procedure

1. Enter `telcset dual_cache_mode primary`.
2. Reboot the server, using the supervisor-level `reboot-app` command. In a cluster, reboot one node at a time.
3. Set the `telc` on all nodes.

Postrequisites

To reenable the IOPS component, delete the `dual_cache_mode telc`, and reboot.

If the throughput component is not enabled when the license is installed

If a Performance Accelerator license is installed, but the throughput component is not enabled, the most likely reason is that the eight-lane connection to the Tachyon Fibre Channel controller has not been successfully tested. For the eight-lane connection to be tested, the server must be completely rebooted, using a full system reboot, and there must be no stale data in NVRAM. If these conditions are met, then the full POST test should run on boot (assuming it has not previously passed).

If the full POST test has not previously passed, and if the test is still not running on boot, check that the license is installed, a full system reboot is being performed, and that there is no stale data in NVRAM. Stale data is cleared from NVRAM by unmounting file systems thoroughly, and using the `nvpages list` command to inspect for stale data.

If the full POST test is running and failing, it might indicate a fault in the server.

The following events are logged by Performance Accelerator:

Event	Description
Performance Accelerator throughput enabled	When Performance Accelerator throughput is enabled, when it was previously disabled.
Performance Accelerator throughput disabled	When Performance Accelerator throughput was enabled but now is not.
Cannot enable Performance Accelerator throughput	When Performance Accelerator is licensed, but POST was not able to run, or it failed to run.

Providing an SSL certificate to the external SMU

Both the server and the SMU are preconfigured with default SSL certificates. These default certificates should provide an acceptable level of security for most users. For added security, this certificate may be replaced with a certificate signed by a *certificate authority* (for example, Verisign).

To request a certificate from a certificate authority (CA):

- Generate a custom private key (optional)
- Generate a certificate signing request (CSR)

By default, all protocols and cipher sties are enabled. However, occasionally a protocol or cipher suite may be no longer secure and the admin can use the **Security Options** page to prevent a browser from communicating with the SMU using that protocol or suite

To disable individual protocols or cipher suites:

- Disable protocols and cipher suites

- ☐ [Generating a custom private key and SSL certificate](#)
- ☐ [Generating a certificate signing request \(CSR\)](#)
- ☐ [Accepting self-signing certificates](#)

Generating a custom private key and SSL certificate

The SMU already contains a default private key from which a CSR can be generated. Default values include:

- Common name (CN) uses the SMU host name, but other values are static (for example: *OU=.*, *O=HDS*, *L=San Jose*, *ST=CA*, *C=US*)
- Valid for 3,650 days (10 years)
- Key length of 2,048 bits

From the SMU CLI, enter `cert-showall.sh` to display these default certificate values.



Note: See the *Hitachi Unified Storage File Module System Access Guide* for directions on how to access the SMU CLI.

To generate a custom private key using other values:

Procedure

1. SSH in to the SMU as the user `manager`, enter `su-` and enter the root password.
2. Enter `cert-gencustom.sh`
3. Enter the requested information as the prompts appear (pressing `Enter` accepts the default).
 - Organizational Unit (OU)
 - Organization (O)
 - Location (L)
 - State (ST)
 - Country (C)
 - Valid Period (in days)
 - Key Size (for example: 1024, 2048 – the key length must be divisible by 64)

After the system confirms the input, it generates a new private key and self-signed certificate.

4. Restart the web server when prompted so that it can start to use the new SSL certificate.
5. Close and restart any browsers that are connected to the SMU. Restarting the browser is required to purge the browser of any previously negotiated SSL session values. When logging into Web Manager after restarting the browser, the new SSL certificate will be provided.
6. To back up the private key and certificate, navigate to **Home > SMU Administration > SMU Backup > Backup**, and save the resulting zip file to a safe and secure location.

The zip file contains a full backup of the SMU's configuration. The `smu.keystore` file within the zip file contains the SMU's private key.

Generating a certificate signing request (CSR)

A certificate signing request (CSR) is a file that contains the encoded information needed to request a certificate from an authority. After generating the CSR, it can be submitted to the authority.

To generate a CSR:

Procedure

1. SSH in to the SMU as user `manager`, enter `su -`, and enter the root password.
2. Enter `cert-gencustom.sh`
3. Copy and paste the CSR that is displayed into the website of the certificate authority.

Alternatively, copy the CSR from the following file on the SMU: `/etc/opt/mercury-papi/ssl/certreq.csr`

Installing certificates

After obtaining the signed certificate from the certificate authority (CA):

Procedure

1. Copy the certificate provided by the CA to the SMU (for example, use the `scp` command to copy the certificate to `/home/manager/signedcert.der.p7`).
2. If necessary, provide the certificate authority's trusted certificate chain as a file (for example, `/home/manager/veritas.pem`). The SMU already includes popular certificate authority trust chains, so this step can typically be skipped



Note: The content of the certificate and trust chain files should only start with `-----BEGIN` and end with `-----END CERTIFICATE-----`.

3. Log in to the SMU as user `manager`, enter `su -`, and enter the root password.
4. If you are using your own private/corporate CA, you will probably need to import that CA certificate.
 - If the root CA certificate and your signed certificate are bundled into a single file (usually a `.p7b` file): Enter `cert-import.sh -p path to` to import your signed certificate and the certificate authority certificate.
 - If the Root CA certificate is in a separate file from your certificate: Enter `cert-importtrustchain.sh -p path to root CA cert file -a`

unique alias to import the certificate authority certificate (usually a `.cer` file) . This is optional and is only required if the java keystore does not already trust the root CA. This might require multiple files or chains, so repeat as necessary. Enter `cert-import.sh -p path to cert file` to import your signed certificate reply (usually a `.p7b` file).
The default SMU SSL certificate is now replaced by your CA-signed certificate.



Note: Any unique alias may be used. If the alias already exists in the SMU's keystore, you will be prompted to replace the old certificate or cancel the import.

5. When prompted to overwrite the existing certificate, enter `y`.
6. Restart the web server when prompted so that it can start to use the new SSL certificate.
7. Close and restart any browsers used to connect to the SMU.
Restarting the browser is required to purge the browser of any previously negotiated SSL session values.
When logging into Web Manager after restarting the browser, the new SSL Certificate is provided.
8. As needed, enter `cert-showall.sh` to display and verify the contents (SSL certificate and trust chain) of the keystore.

Recreating the default SMU certificate

If there are problems when trying to create/import an SSL certificate, the SMU's default certificate may be recreated.

To recreate the default certificate:

Procedure

1. Log in to the SMU as the user `manager`, enter `su -`, and enter the root password.
2. Enter `cert-gendefault.sh`.
3. When prompted to overwrite the existing certificate, enter `y`.
4. Restart the web server when prompted so that it can use the new SSL certificate.
5. Close and restart any browsers that are connected to the SMU.
Restarting the browser is required to purge the browser of any previously negotiated SSL session values. When logging into Web Manager after restarting the browser, the new SSL certificate will be provided.

Accepting self-signing certificates

If a self-signed certificate has been installed, you receive a security alert when you first access the Web Manager over a secure connection.

Procedure

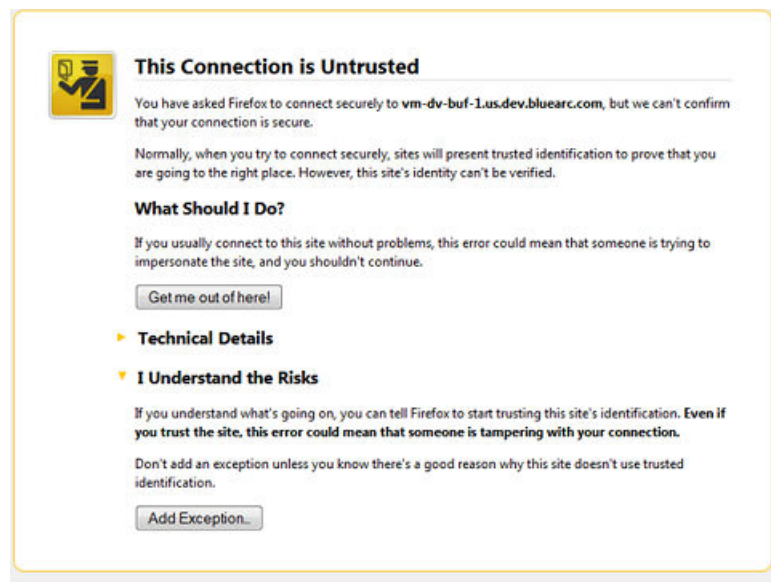
1. If a self-signed certificate has been installed, you receive a security alert when you first access the Web Manager over a secure connection. Although you can click **Yes** to proceed, the alert reappears when you next run the Web Manager. To suppress the alert, you must opt to trust the certifying authority

- **For Internet Explorer:**

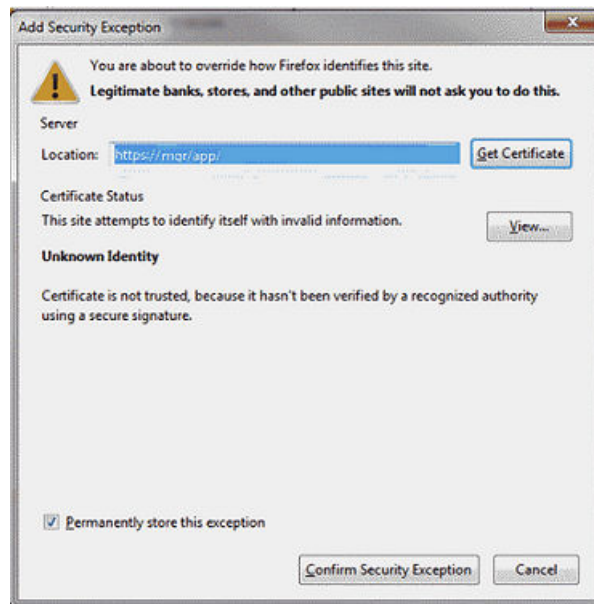
- a. From the **Security Alert** dialog, click **View Certificate** to display the certificate.
- b. Click **Install Certificate**, then follow the on-screen instructions to install the certificate in the Trusted Root Certification Authorities store.

- **For Firefox:**

- a. Firefox-based browsers display an alert message.



- b. Accepting the certificate permanently suppresses the alert in future sessions.



Providing an SSL certificate to the embedded SMU

You can use HNAS console commands to restrict which versions of SSL/TLS and cipher suites may be used to comply with your IT security policies. And you can also replace the default "self-signed" certificate with your own Certificate Authority (CA) signed certificate.

To restrict which versions of SSL/TLS and cipher suites may be used and to provide your own Certificate Authority (CA) signed certificates for use in the HNAS internal SMU:

- Configure cipher suites
- Configure the SSL/TLS version
- Obtain and import a CA-signed certificate

- ☐ [Configuring cipher suites](#)
- ☐ [Configuring the SSL/TLS version](#)
- ☐ [Obtaining and importing a CA-signed certificate](#)

Configuring cipher suites

You can restrict which cipher suites may be used to comply with your security policies.

Use HNAS console commands to configure cipher suites to disable cipher suites you do not wish to use.

Procedure

1. To list the enabled cipher suites, enter:

```
$ tls-cipher-suite-list
```



Note: See the *Hitachi Unified Storage File Module System Access Guide* for directions on how to access the HNAS server CLI.

The enabled and disabled cipher suites are shown.

2. To list specific cipher suites, enter:

```
$ tls-cipher-suite-list EXP-RC4-MD5
$ tls-cipher-suite-list EXP-RC4-MD5:  enabled
```

The `tls-cipher-suite-list` command lists all known cipher suites and shows whether each is enabled or disabled.

3. To disable an enabled cipher suite, enter:

```
$ tls-cipher-suite-disable --confirm EXP-RC4_MD5
$ tls-cipher-suite-list EXP-RC4-MD5
EXP-RC4-MD5 : disabled
```



Note: The `--confirm` option must be included to commit changes and restart the HTTPS server.

4. To enable a disabled cipher suite, enter:

```
$ tls-cipher-suite-enable --confirm EXP-RC4_MD5
$ tls-cipher-suite-list EXP-RC4-MD5
EXP-RC4-MD5 : enabled
```

5. To reset the cipher suites to the defaults, enter:

```
$ tls-cipher-suite-default --confirm
```

Result

When the SSL configuration is changed, or a custom certificate is installed or removed, the HTTPS management server is automatically restarted to ensure that all current and future connections make use of the certificate, and the enabled versions and ciphers. An incorrect configuration can cause the SMU

to be unable to communicate with the HTTPS management server. Verify that the SMU can still communicate after the settings have been changed.

Configuring the SSL/TLS version

You can restrict which versions of SSL/TLS may be used to comply with your security policies.

Use the following commands to configure the SSL/TLS version and restrict which versions of SSL/TLS may be used.

Procedure

1. List the enabled SSL/TLS versions:

```
$ tls-version-list
SSLv2      : disabled
SSLv3      : disabled
TLSv1      : enabled
TLSv1.1    : enabled
TLSv1.2    : enabled
```

2. Set the enabled SSL/TLS versions. The SMU supports TLSv1.2, so it is recommended that you use this version.

```
$ tls-version-set --tls1.1 --tls1.2 --confirm
```



Note: You should not enable SSLv2, because it is not secure.

3. Set the enabled SSL/TLS versions to the default. The default versions are TLS1.0, TLS1.1 and TLS1.2 enabled, and SSL2 and SSL3 disabled.



Note: These default values are currently safe, but this may change as vulnerabilities are found in different SSL/TLS versions.

```
$ tls-version-set --default --confirm
```

Result

When the SSL configuration is changed, or a custom certificate is installed or removed, the HTTPS management server is automatically restarted to ensure that all current and future connections make use of the certificate, and the enabled versions and ciphers. An incorrect configuration can cause the the SMU to be unable to communicate with the HTTPS management server. Verify that the SMU can still communicate after the settings have been changed.

Obtaining and importing a CA-signed certificate

You may provide your own Certificate Authority (CA) signed certificates, instead of the default "self-signed" certificate.

Use these steps to obtain and import a CA-signed certificate into the server.

Prerequisites

Supported encoding of the certificates are PEM or DER.

The trust chain certificates must be in X.509 format.

The signed certificate must be in X.509 format or a PKCS #7 bundle that includes the trust chain certificates.

Procedure

1. Create a new certificate. Customize the server's private key to set the required validity period and correct location information.

```
$ tls-certificate-create-custom --confirm
```

2. Generate a CSR (Certificate Signing Request) and send it to the chosen CA.

```
$ tls-certificate-generate-csr
```



Note: The CA will check the sender's identity. This may take some time.

3. Depending on what you are provided, perform the appropriate steps:
 - If you are given a single X.509 signed certificate and multiple X.509 trust chains:

1. Import each certificate of the trust chain provided.

```
$ tls-certificate-create-custom --confirm --path  
tcl.cer -alias tcl
```

```
$ tls-certificate-create-custom --confirm --path  
tcn.cer --alias tcn
```

2. Import the signed certificate.

```
$ tls-certificate-import-signed --confirm --path  
signed.cer
```

- If you are given a single PKCS #7 certificate bundle:
Depending on the format of the trust chain and signed certificate, you may import them both at once.

```
$ tls-certificate-import-signed --confirm --path  
signed_and_trust_chain
```

Result

When the SSL configuration is changed, or a custom certificate is installed or removed, the HTTPS management server is automatically restarted to ensure that all current and future connections make use of the certificate, and the enabled versions and ciphers. An incorrect configuration can cause the the SMU to be unable to communicate with the HTTPS management server. Verify that the SMU can still communicate after the settings have been changed.

Using HNAS multi-tenancy

The HNAS multi-tenancy feature provides HNAS application service providers (ASPs) with another configuration mode option in addition to the standalone HNAS individual EVS security feature. Both options provide support for multiple file serving Enterprise Virtual Servers (EVSs) on a single HNAS host or multiple hosts. However, the multi-tenancy option extends the functionality of the stand-alone option and provides additional security and configuration enhancements.

- ☐ [Using HNAS multi-tenancy](#)
- ☐ [Understanding multi-tenancy](#)
- ☐ [Understanding HNAS multi-tenancy benefits](#)
- ☐ [How multi-tenancy mode differs from stand-alone mode](#)
- ☐ [How multi-tenancy differs from per-EVS security](#)
- ☐ [Multi-tenancy requirements](#)
- ☐ [Disabling HNAS multi-tenancy](#)
- ☐ [Managing multi-tenancy](#)
- ☐ [Overlapping IP address support for HNAS multi-tenancy](#)
- ☐ [Understanding routing by EVS](#)
- ☐ [Configuring routes per EVS](#)
- ☐ [Understanding EVS crosstalk checking](#)
- ☐ [Multi-tenancy-aware protocols](#)

Using HNAS multi-tenancy

The HNAS multi-tenancy feature provides HNAS application service providers (ASPs) with another configuration mode option in addition to the standalone HNAS individual EVS security feature. Both options provide support for multiple file serving Enterprise Virtual Servers (EVSs) on a single HNAS host or multiple hosts. However, the multi-tenancy option extends the functionality of the stand-alone option and provides additional security and configuration enhancements.

Understanding multi-tenancy

Multi-tenant architecture provides companies, such as application service providers (ASPs), the ability to support more than one customers' services on a single server, but still keep them logically separate.

In an HNAS server implementation, this architecture is sometimes called *real/ EVS separation*.



Note: The ASP has the responsibility of managing the storage, file systems, shares, and exports to which each tenant has access.

HNAS multi-tenancy configuration mode provides enhancements to the previous stand-alone mode in the following ways:

- Supports tenant configurations in logically separate serving environments on a single physical server or cluster.
- Extends HNAS individual security mode to provide true separation by maintaining per-EVS variables and connection states.
- Supports serving environments for tenants with single or multiple EVSs, configured separately and possibly sharing file serving interfaces.
- Provides per-EVS IP routing and networking settings to support duplicate or overlapping server IP addresses. Includes support for both IPv4 and IPv6.
- Helps detect and prevent EVS crosstalk that can occur when duplicate IP ranges are used. EVS crosstalk can lead to server unresponsiveness.
- Provides CLI EVS context usability improvements.

See the following example of an HNAS multi-tenancy setup:

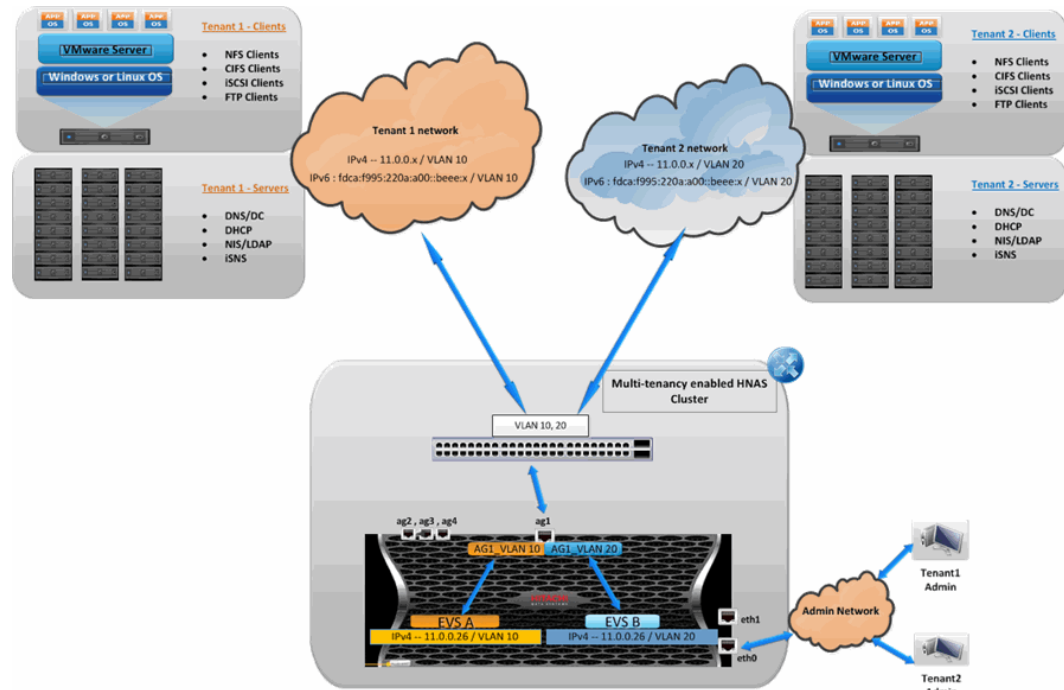


Figure A-1 HNAS multi-tenancy

Understanding HNAS multi-tenancy benefits

Using HNAS multi-tenancy can help you avoid some of the challenges faced with traditional multi-tenant environments.

Commonly, HNAS customers who are ASPs (Internet services providers and managed services providers) sell their services to their customers. Their customers are the tenants in a multi-tenant environment. The ASPs cannot force their tenants into a specific subnet, which means that the ASPs run into issues when some tenants use the same network address scheme.

In the past, this situation caused overlapping IP addresses and networks on the HNAS EVSs. The IP routing and networking settings were global on an HNAS server--per-EVS settings were unsupported. The HNAS multi-tenancy feature allows you to set up all the different tenant networks as VLANs and then allocate them to the specific EVSs. These networks may have the same IP subnet but may be different gateways in their VLAN-segregated networks.

Other common challenges that HNAS multi-tenancy addresses include the following:

- Tenant's configurations could contain the same names but identify different things to those tenants. For example, the names could identify NIS domain, Windows domain, or virus scanners.
- Tenant's networks could have the same address range. They could be in the same or different subnets from each other.

- Tenant's networks could be using the same IP address for something, but the server is really a different server. For example, it's common practice to use the first or last IP address in a given subnet to be the router for that network, so the same IP address could be referring to many different routers for different customers.
- Looking up a name from a given IP address may give different names for different tenants.
- Looking up an IP address from a name may give different IP addresses for different tenants. Even if they resolve to the same IP address, they may be completely different hosts.

How multi-tenancy mode differs from stand-alone mode

The HNAS multi-tenancy mode option provides additional security and configuration enhancements.

Capabilities	Configuration mode	
	Stand-alone	Multi-tenancy
Multiple EVSs per HNAS	x	x
Logically separate serving environments on a single HNAS or cluster		x
Combining multiple EVSs into one EVS		x
Per-EVS security with global namespaces	x	x
Legacy VLANs (deprecated)	x	
VLAN-interface	x	x
Duplicate or overlapping IP address support		x
EVS crosstalk checking		x
Per-EVS routing	x	x
Multi-tenancy-aware protocols		x

How multi-tenancy differs from per-EVS security

Both HNAS configuration modes provide per-EVS security.

The following table shows a comparison of per-EVS security support for both modes:

Per-EVS security capabilities	Configuration mode	
	Stand-alone	Multi-tenancy
Nameservice	x	x
DNS configuration	x	x
Routes		x
Netbios supported	x	
Separate IP address spaces		x
Others		x

Multi-tenancy requirements

Requirements for enabling and using multi-tenancy mode.

In order to enable and use multi-tenancy mode, the following requirements must be met:

- The per-EVS security license must be installed.
- All EVSs present on the NAS server or cluster must be configured with individual security settings.
If you want to convert an NAS server to use multi-tenancy, contact technical support.
- An EVS may use a VLAN interface, or an aggregation interface, but VLANs configured with the `vlan` command are not supported. VLAN interfaces are configured using the `vlan-interfaces` command.
Scripts are available to convert the VLANs to VLAN interfaces, but these scripts should not be used without technical support guidance.



Note: When multiple EVSs are used for a tenant, no policing of aggregation interface or VLAN interface usage is performed. The NAS server administrator must ensure that EVSs for different tenants do not use the same network interface. If EVSs for different tenants do use the same interface, traffic for one tenant could be seen on a different tenant's network (this situation would be no different than if a switch was configured incorrectly).

- No cluster name space (CNS) may be configured (an EVS name space is supported).
- Active Directory Server (ADS) entries must be used instead of NT domains.



Note: When multi-tenancy is enabled, NetBIOS is disabled and NT4 domains cannot be used.

- For clusters, all nodes must be running a version of software that supports multi-tenancy
- When enabling multi-tenancy mode for a cluster, all cluster nodes must be online.

Disabling HNAS multi-tenancy

In most cases, the **multi-tenancy-disable** command will only be run when performing a major reconfiguration or decommissioning of an existing system. See the **multi-tenancy-disable** command in the CLI man pages for further details.



Note: All EVSs must be deleted prior to issuing this command

Managing multi-tenancy

Managing multi-tenancy on the NAS server and EVSs.

Multi-tenancy is an operational mode for the NAS server, which allows the NAS server to support multiple tenants, each with at least one EVS. Multi-tenancy supports tenant configurations in logically separate serving environments on a single physical server or cluster. Each EVS has an individual configuration and is managed as an individual unit, without sharing a global configuration. In stand-alone (single tenant) mode, the NAS server has a global configuration which can be shared by EVSs.

Multi-tenancy management interfaces

Currently Web Manager does not include support for enabling and disabling multi-tenancy; to use this feature, you must use CLI commands. See the following CLI man pages for detailed information on configuring and using multi-tenancy:

The following commands are used to manage multi-tenancy at the NAS server level:

- **multi-tenancy-disable**
- **multi-tenancy-enable**
- **multi-tenancy-show**

For more information about these commands, refer to the *Command Line Reference*. For an overview of multi-tenancy related commands and other related information, see the **multi-tenancy** command in the *Command Line Reference*.

Viewing HNAS multi-tenancy status

Use the **multi-tenancy-show** command to view the status of the multi-tenant environment.

Command example:

```
server:$ multi-tenancy-show
Multi-tenancy is disabled.
server:$
```

Considerations for enabling HNAS multi-tenancy

In most cases, this command will only be run when first configuring a new system.

Enabling the multi-tenant environment causes a temporary loss of service to all EVSs while they are enabled for use.



Note: All connections to the server are disconnected during the enabling process. This includes the connection that may be used to execute the enabling command.

HNAS multi-tenancy limits

Multi-tenancy can only be enabled if all cluster nodes are multi-tenant capable, and are all online. It also requires an EVS Security license.

Multi-tenancy cannot be enabled:

- if any EVS is configured with global security settings. Resolve this by setting the EVSs to individual security by issuing the **evs-security** command and copying the required configuration.
- if a deprecated VLAN subnet is still configured in the system, a VLAN interface must be used. Remove VLAN subnets by using the **vlan remove-all** command. Create appropriate VLAN interfaces using the **vlan-interface-create** command.
- If a cluster namespace is configured, issue the **namespace-delete** command before enabling multi-tenancy.

Enabling HNAS multi-tenancy

To enable HNAS multi-tenancy, issue the **multi-tenancy-enable** command.

Enabling the multi-tenant environment will cause a temporary loss of service to all EVSs while they are enabled for use.

Command Example:

```
server:$ multi-tenancy-enable
Warning: Enabling multi-tenancy significantly affects the
configuration of the HNAS.
Have you read and understood the multi-tenancy man page? (Y/N) [N]:
Y
Have you read and understood the multi-tenancy-enable man page?
```

```
(Y/N) [N]:
Y
Do you understand that once enabled, multi-tenancy cannot be
disabled until all file serving EVSs have been deleted?(Y/N) [N]:
Y
Warning: All active connections, including any remote console
sessions, will be disconnected to allow the network service to
support multi-tenancy.
Do you want to proceed?(Y/N) [N]:
Y
Enabling multi-tenancy.
Operation successful.
server:$
```

Managing multi-tenancy on the NAS server

The administration of the NAS server is performed by the NAS server administrator of the application service provider (ASP), who has access to the administrative EVS. The NAS server administrator of the ASP is responsible for managing the storage, file systems, shares, and exports for the tenant.

NAS server multi-tenancy is enabled and managed through CLI commands. Management of shares, exports, file systems, and storage pools/spans, replication, migration, and any other functions that are not EVS-specific performed in the same way regardless of whether the NAS server is operating in multi-tenant mode or in single tenant (stand-alone) mode.

The following commands are used to manage multi-tenancy at the NAS server level:

- **multi-tenancy-disable**
- **multi-tenancy-enable**
- **multi-tenancy-show**

For more information about these commands, refer to the *Command Line Reference*. For an overview of multi-tenancy related commands and other related information, see the **multi-tenancy** command in the *Command Line Reference*.

Managing multi-tenancy for an EVS

Multi-tenancy extends HNAS individual security mode to provide true separation by maintaining per EVS variables and connection states. Each EVS has its own complete and separate environment which is configured as required (for example, network interfaces, routing, and security aspects are set on a per-EVS basis, instead of a global configuration).

When managing an EVS through the command line interface, the individual EVS context (label) is displayed on the command line prompt, and commands will affect only that EVS. Use the **console-context** command to change the current context.

The following commands are used to manage the EVS when multi-tenancy is enabled:

- `evsipaddr`
- `set-for-evs`
- `routing-by-evs`
- `routing-by-evs-disable`
- `routing-by-evs-enable`
- `routing-by-evs-show`

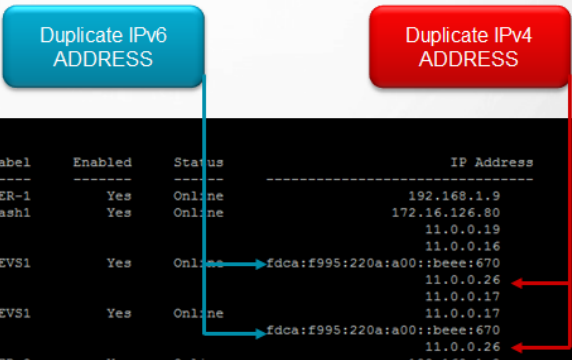
For more information about these commands, refer to the *Command Line Reference*. For an overview of multi-tenancy related commands and other related information, see the `multi-tenancy` command in the *Command Line Reference*.

Overlapping IP address support for HNAS multi-tenancy

HNAS configuration

Enabling multi-tenancy enables support for configuring duplicate/overlapping server IP addresses and IP subnets governed by the following rules:

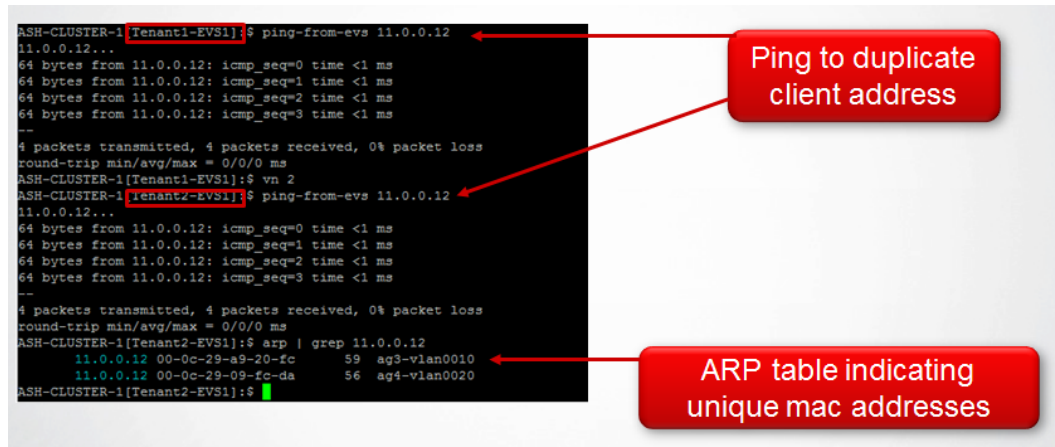
- The same IP address is not used by the same EVS more than once.
- The same IP address is not used by the same interface more than once where an interface can be an aggregate or an aggregate VLAN



```
ASH-CLUSTER-1:$ evs list
```

Node	EVS ID	Type	Label	Enabled	Status	IP Address	Port
1		Cluster	ASH-CLUSTER-1	Yes	Online	192.168.1.9	eth1
1	0	Admin	merc-ash1	Yes	Online	172.16.126.80	eth0
						11.0.0.19	ag3-vlan0010
						11.0.0.16	ag4-vlan0020
1	1	Service	Tenant1-EVS1	Yes	Online	fdca:f995:220a:a00::beee:670	ag3-vlan0010
						11.0.0.26	ag3-vlan0010
						11.0.0.17	ag3-vlan0010
1	2	Service	Tenant2-EVS1	Yes	Online	11.0.0.17	ag4-vlan0020
						fdca:f995:220a:a00::beee:670	ag4-vlan0020
						11.0.0.26	ag4-vlan0020
2		Cluster	ASH-CLUSTER-2	Yes	Online	192.168.1.8	eth1

Discovering duplicate client addresses



Understanding routing by EVS

Routing by EVS restricts the choice of source addresses available to the routing engine to those associated with the source EVS. Routing by EVS is always enabled in multi-tenancy mode. Routing by EVS can also be enabled when not in multi-tenancy mode.

Some subsystems already use the current EVS to influence routing decisions. With routing by EVS enabled, many subsystems, such as DNS, which normally would not use the EVS to influence routing decisions, now would use routing by EVS. If routing by EVS is to be enabled in non-multi-tenant mode, it is necessary to use the **routing-by-evs-enable** command. See the CLI reference for **routing-by-evs** commands:

- **routing-by-evs-enable**
- **routing-by-evs-disable**
- **routing-by-evs-show**

Configuring routes per EVS

Multi-tenancy causes the routing engine to keep routes by EVS, so it is necessary to maintain different sets of routes for each EVS. Gateway, network and host routes (IPv4 and IPv6) are configured per EVS when multi-tenancy is enabled using the following commands: **route-gateway-add**, **route-net-add**, and **route-host-add**.

Configuration is done with the EVS in context. Prefix lengths are accepted for IPv4 and IPv6 network addresses.

Command examples:

```
hnas:$ evs-select 1
hnas[EVS01]:$ route-gateway-add fdca:f995:220a:a00::1
Route cache flushed.
```

```

hnas[EVS01]:$ evs-select 2
hnas[EVS02]:$ route-net-add 10.2.0.0/16 -g 10.1.2.3 -m 9000
Route cache flushed.
hnas[EVS02]:$ evs-select 3
hnas[EVS03]:$ route-host-add 10.1.2.3 -g 10.1.3.4
Route cache flushed.

```

The route command is display only when multi-tenancy is enabled. The route command displays routes for the EVS in context.

```

[EVS01]:$ route
Routes for EVS 1:
Destination          Gateway              MTU    Flags
::/0                 fdca:f995:220a:a00::1  Default G
::/0                 fe80::208:e3ff:feff:fc28  1500  GD    via eth0
::/0                 fe80::208:e3ff:feff:fc28  1500  GD    via ag1

[EVS02]:$ route
Routes for EVS 2:
Destination          Gateway              MTU    Flags
10.2.0.0/16          10.1.2.3             9000
::/0                 fe80::208:e3ff:feff:fc28  1500  GD    via eth0
::/0                 fe80::208:e3ff:feff:fc28  1500  GD    via ag1

```

The route command will redirect the user to the route commands for configuration when multi-tenancy is enabled.

```

[Tenant2-EVS2]:$ route add gateway
route: as multi-tenancy is enabled, use route-gateway-add

```

```

[Tenant2-EVS2]:$ route add host
route: as multi-tenancy is enabled, use route-host-add

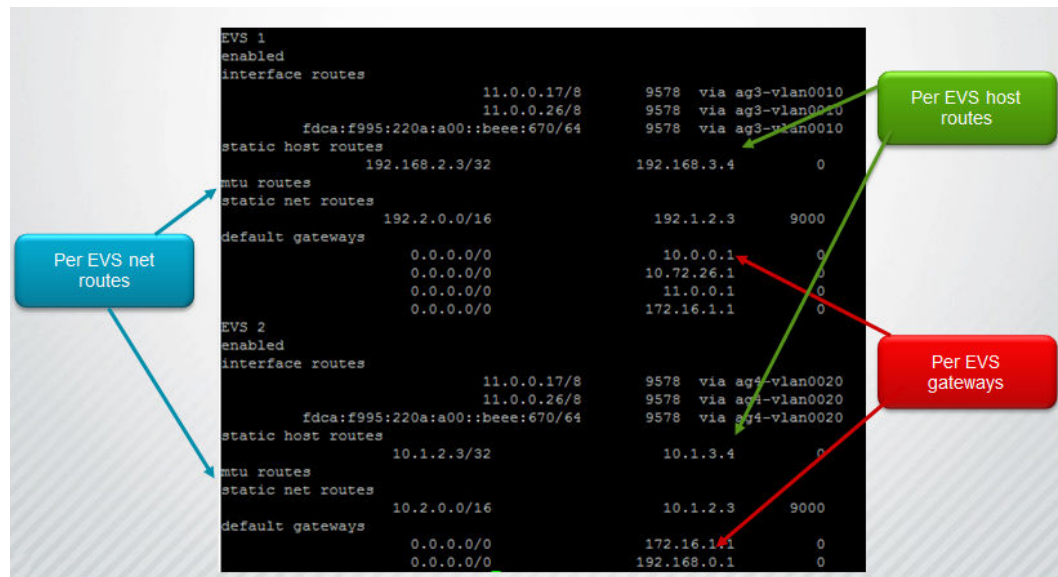
```

```

[Tenant2-EVS2]:$ route add net
route: as multi-tenancy is enabled, use route-net-add

```

The commands **router-dump-by-evs** and **test-route-by-evs** can be used to diagnose networking problems, where routes are configured per EVS.



Understanding EVS crosstalk checking

The HNAS platforms support detection and prevention of EVS crosstalk. Crosstalk can cause the server to fail to respond. Crosstalk checking is especially important when duplicate IP ranges are being used.

Enabling multi-tenancy automatically enables crosstalk checking. The checks within the code ensure reliable packet delivery.

Multi-tenancy-aware protocols

The HNAS multi-tenancy mode feature recognizes and uses certain protocols. This mode extends the previous stand alone mode protocol support.

HNAS multi-tenancy supports the following protocols:

- CIFS/SMB
- NFS
- FTP
- iSCSI

Consider the following characteristics of how the protocols are supported:

- Incoming requests and outgoing responses are made on a per-EVS basis.
- Protocol stack crosschecks the IP to VNODE ID mapping against the EVS ID passed by the network stack.
- Configuration and connection states are maintained on a per-EVS basis.

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas

+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa

+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific

+852 3189 7900
hds.marketing.apac@hds.com



MK-92HNAS010-08