



Overland
Storage

SnapScale™

Administrator's Guide

For a Clustered Network Running
RAINcloudOS™ Version 4.0



December 2013
10400455-001



Dec 2013 ©Overland Storage, Inc. All rights reserved.

Overland®, Overland Data®, Overland Storage®, ARCVault®, DynamicRAID®, LibraryPro®, LoaderXpress®, Multi-SitePAC®, NEO®, NEO Series®, PowerLoader®, Protection OS®, REO®, REO 4000®, REO Series®, Snap Appliance®, Snap Care® (EU only), SnapServer®, StorAssure®, Ultamus®, VR2®, and XchangeNOW® are registered trademarks of Overland Storage, Inc.

GuardianOS™, RAINcloud™, RapidRebuild™, SnapDisk™, SnapEDR™, Snap Enterprise Data Replicator™, SnapExpansion™, SnapSAN™, SnapScale™, SnapServer DX Series™, SnapServer Manager™, and SnapWrite™ are trademarks of Overland Storage, Inc.

All other brand names or trademarks are the property of their respective owners.

The names of companies and individuals used in examples are fictitious and intended to illustrate the use of the software. Any resemblance to actual companies or individuals, whether past or present, is coincidental.

PROPRIETARY NOTICE

All information contained in or disclosed by this document is considered proprietary by Overland Storage. By accepting this material the recipient agrees that this material and the information contained therein are held in confidence and in trust and will not be used, reproduced in whole or in part, nor its contents revealed to others, except to meet the purpose for which it was delivered. It is understood that no right is conveyed to reproduce or have reproduced any item herein disclosed without express permission from Overland Storage.

Overland Storage provides this manual as is, without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Overland Storage may make improvements or changes in the product(s) or programs described in this manual at any time. These changes will be incorporated in new editions of this publication.

Overland Storage assumes no responsibility for the accuracy, completeness, sufficiency, or usefulness of this manual, nor for any problem that might arise from the use of the information in this manual.

FW 4.0.061

Overland Storage, Inc.
9112 Spectrum Center Blvd.
San Diego, CA 92123
U.S.A.

Tel: 1.877.654.3429 (toll-free U.S.)
Tel: +1.858.571.5555, Option 5 (International)
Fax: +1.858.571.0982 (general)
Fax: +1.858.571.3664 (sales)
www.overlandstorage.com

Audience and Purpose

This guide is intended for system and network administrators charged with installing and maintaining a SnapScale cluster running RAINcloudOS 4.0 on their network. It provides information on the installation, configuration, security, and maintenance of the SnapScale cluster and nodes.

Product Documentation

SnapScale product documentation and additional literature are available online, along with the latest release of the RAINcloudOS 4.0 software.

Point your browser to:

<http://docs.overlandstorage.com/snapscale>

Follow the appropriate link on that page to download the **latest** software file or document.

For additional assistance, search at <http://support.overlandstorage.com>.

Overland Technical Support

For help configuring and using your SnapScale cluster, email our technical support staff at:

techsupport@overlandstorage.com.

You can get additional technical support information on the [Contact Us](#) web page at:

<http://docs.overlandstorage.com/support>

For a complete list of support times based on your type of coverage, visit our website at:

<http://docs.overlandstorage.com/care>

Software Updates

The latest release of the RAINcloudOS software can be obtained from the Downloads and Resources (SnapScale Solutions) page at the Overland Storage website:

<http://docs.overlandstorage.com/snapscale>

Follow the appropriate instructions to download the **latest** software file.

For additional assistance, search at <http://support.overlandstorage.com/>.

Conventions

This document exercises several alerts and typographical conventions.

Alerts

Convention	Description & Usage
 IMPORTANT	An <i>Important</i> note is a type of note that provides information essential to the completion of a task or that can impact the product and its function.
 CAUTION	A <i>Caution</i> contains information that the user needs to know to avoid damaging or permanently deleting data or causing physical damage to the hardware or system.
 WARNING	A <i>Warning</i> contains information concerning personal safety. Failure to follow directions in the warning could result in bodily harm or death.
AVERTISSEMENT	Un Canadien <i>avertissement</i> comme celui-ci contient des informations relatives à la sécurité personnelle. Ignorer les instructions dans l'avertissement peut entraîner des lésions corporelles ou la mort.

Typographical Conventions

Convention	Description & Usage
Button_name	Words in this special boldface font indicate the names of command buttons found in the Web Management Interface.
Ctrl-Alt-r	This type of format details the keys you press simultaneously. In this example, hold down the Ctrl and Alt keys and press the r key.
NOTE	A Note indicates neutral or positive information that emphasizes or supplements important points of the main text. A note supplies information that may apply only in special cases, for example, memory limitations or details that apply to specific program versions.
Menu Flow Indicator (>)	Words with a greater than sign between them indicate the flow of actions to accomplish a task. For example, Setup > Passwords > User indicates that you should press the Setup button, then the Passwords button, and finally the User button to accomplish a task.
<i>Courier Italic</i>	A variable for which you must substitute a value.
Courier Bold	Commands you enter in a command-line interface (CLI).

Information contained in this guide has been reviewed for accuracy, but not for product warranty because of the various environments, operating systems, or settings involved. Information and specifications may change without notice.

Japanese Voluntary Control Council for Interference (VCCI)

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI— A

(Translation: This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case, the user may be required to take corrective actions.)

Preface

Chapter 1 - Overview

SnapScale Conventions	1-1
SnapScale Node Requirements	1-2
RAINcloudOS Specifications	1-3
RAINcloudOS 4.0 Features	1-4
Client and Storage Networks	1-5
Node/Switch Cabling Example	1-5
Node Port Configurations	1-6
X2 Node Configurations	1-6
X4 Node Configurations	1-7

Chapter 2 - Setup and Configuration

Connecting for the First Time	2-1
Connect Using the Node Name	2-1
Connect Using SSM	2-2
Create a New SnapScale Cluster (via Wizard)	2-3
Step 1 – Select SnapScale Nodes	2-3
Step 2 – Client Network Configuration Overview	2-5
Step 3 – Choose Client Network Static TCP/IP Settings	2-5
Step 4 – Configure Node Static IP Addresses	2-6
Step 5 – Basic SnapScale Properties	2-7
Step 6 – Set Date and Time	2-8
Step 7 – Summary Verification & Cluster Creation	2-9
Join an Existing SnapScale Cluster (via Wizard)	2-12
Web Management Interface	2-13
Alert Messages	2-15
Site Map	2-16
Contact Information	2-17

Chapter 3 - SnapScale Settings

SnapScale Properties	3-2
Date/Time	3-3
Configure Date and Time Settings Manually	3-3
Configure Date and Time Settings for Automatic Synchronization	3-4
SSH	3-5
UPS	3-5
Edit UPS Properties	3-6
Procedure to Configure UPS Protection	3-7
Add Network UPS Device	3-7

Change Network UPS Device	3-8
Delete Network UPS Device	3-8

Chapter 4 - Network Access

View Network Information	4-2
Client Network Information	4-2
Storage Network Information	4-4
TCP/IP Networking	4-5
Bonding Options	4-6
Guidelines in TCP/IP Configuration	4-7
Edit Storage Network Properties	4-8
Utility IP Address	4-9
Windows/SMB Networking	4-10
Support for Windows/SMB Networking	4-11
Support for Windows Network Authentication	4-12
Connect from a Windows Client	4-13
Connect a Mac OS X Client Using SMB	4-13
Configure Windows/SMB Networking	4-13
NFS Access	4-16
Support for NFS	4-17
NFS Share Mounting	4-17
NIS Domains	4-18
Guidelines for Configuring NIS	4-18
FTP/FTPS Access	4-19
Supported FTP Clients	4-19
SNMP Configuration	4-20
Default Traps	4-21
Supported Network Manager Applications and MIBs	4-22
Configure SNMP	4-22
Web Access	4-23
Configuring HTTP/HTTPS	4-23
Using Web Root to Configure the SnapScale as a Simple Web Server	4-24
iSNS Configuration	4-27

Chapter 5 - Storage Options

Peer Sets	5-1
Peer Sets and Recovery	5-2
Peer Set Utilization	5-4
Peer Set Basics	5-4
Peer Sets Page	5-5
Spare Disks Page	5-6
Spare Distributor	5-7
Data Balancer	5-9
Volumes	5-11
Volume Overview	5-11
Creating Volumes	5-12
Edit Volume Properties	5-14
Deleting Volumes	5-14
Quotas	5-15
Default Quotas	5-16
Quotas for Volume Page	5-17

Add Quota Wizard	5-18
Editing or Removing Quotas	5-20
Snapshots	5-21
Snapshots Overview	5-22
Creating Snapshots	5-22
Adjusting Snapshot Space	5-25
Accessing Snapshots	5-26
Scheduling Snapshots	5-26
Edit Snapshot Properties	5-27
iSCSI Disks	5-27
Configuring iSCSI Initiators	5-27
iSCSI Configuration on the SnapScale	5-27
Create iSCSI Disks	5-29
Edit iSCSI Disk Properties	5-31
Delete an iSCSI Disk	5-32
Configuring VSS/VDS for iSCSI Disks	5-32
Nodes	5-35
Nodes Overview	5-35
Edit Node Properties	5-36
Node Drives	5-36
Adding Nodes	5-37
Removing Nodes	5-41
Node Identification	5-42
Disks	5-43
Replacing Drives	5-43
Adding Drives	5-44

Chapter 6 - Security Options

Overview	6-1
Guidelines for Local Authentication	6-2
User and Group ID Assignments	6-3
Security Guides	6-3
Windows Active Directory Security Guide	6-4
Entire Volume Security Guide	6-5
Folder on Volume Security Guide	6-5
Shares	6-5
Share Security Overview	6-6
Create Shares	6-6
Edit Share Properties	6-8
Delete Shares	6-9
Configuring Share Access	6-9
Local Users	6-16
Create a User	6-16
Edit User Properties	6-17
Local User Password Policies	6-19
Assign User to Group	6-20
Delete Local User	6-21
Local Groups	6-22
Create New Group	6-22
Edit Group Properties	6-23
Specify Users in Group	6-24

Delete Group	6-24
Security Models	6-25
Managing Volume Security Models	6-25
ID Mapping	6-25
Add Mapping	6-26
Change Mapping	6-29
Auto Mapping	6-31
Remove Mappings	6-33
Remove Missing ID Mappings	6-36
Filesystem Updates	6-37
Home Directories	6-37
Configure Home Directories	6-38

Chapter 7 - System Monitoring

System Status	7-2
SnapScale Status	7-2
Activity	7-3
Active Users	7-3
Open Files	7-4
Network Monitor	7-4
Event Log	7-7
Filter the Log	7-7
Protocol Manager	7-8
SnapScale Settings	7-9
Tape	7-11

Chapter 8 - Maintenance

Shutdown and Restart	8-2
Manually Powering Nodes On and Off	8-2
Data Import	8-2
Setting Up a Data Import Job	8-4
Stopping an Import Job	8-6
Recreating an Import Job	8-6
Preserving Permissions	8-6
OS Update	8-7
Update the RAINcloudOS Software	8-8
Software Update Notification	8-8
Configuring Update Notification	8-9
Manually Checking for Updates	8-9
Support	8-9
Phone Home Support	8-10
Registering Your Cluster	8-12
Tools	8-13
Email Notification	8-14
Host File Editor	8-15
Delete SnapScale Cluster	8-16

Chapter 9 - Misc. Options

Home Pages – Web/Admin	9-1
Web Home	9-2
Administration	9-3

SnapExtensions	9-5
Snap EDR	9-5
Snap Finder	9-6
Edit Snap Finder Properties	9-7
Change Password	9-8
Changing Your Password	9-8
Management Interface Settings	9-9

Appendix A - Backup Solutions

Backup and Replication Solutions	A-1
Snap Enterprise Data Replicator	A-1
Snap EDR Usage	A-2
Configuring Snap EDR for RAINcloudOS	A-2
Scheduling Jobs in Snap EDR	A-3
Backup via SMB or NFS	A-3
Backup via Agent or Media Server	A-3
Utility IP Address	A-3

Appendix B - Security and Access

Security Model Rules	B-1
Security Model Management	B-2
Special Share Options	B-2
Hiding Shares	B-2
Where to Place Shares	B-3
File and Share Access	B-3
Cumulative Share Permissions	B-3
Snapshot Shares and On Demand File Recovery	B-3
Creating a Snapshot Share	B-3
File-level Security	B-4
Security Personalities and Security Models	B-4
Windows ACLs	B-4

Appendix C - RAINcloudOS Ports

Appendix D - Troubleshooting

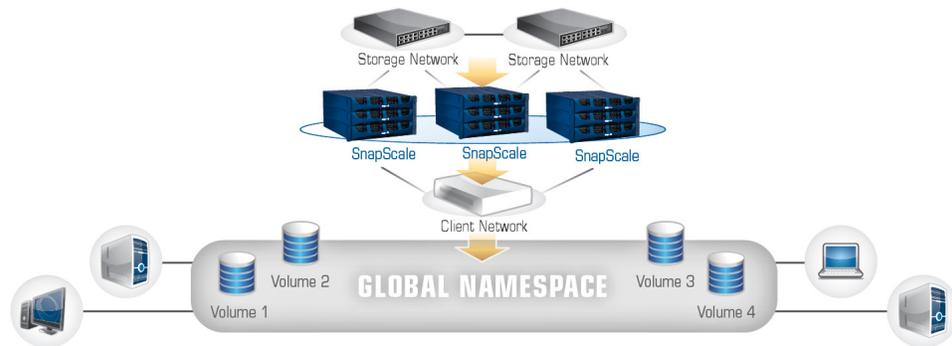
LED Indicators	D-1
SnapScale X2 Node LEDs	D-1
SnapScale X4 Node LEDs	D-3
Network Reset	D-3

Master Glossary & Acronym List

Index

SnapScale is a flexible, scalable, low-maintenance network-attached storage cluster composed of a redundant array of independent nodes running RAINcloudOS. This guide applies to SnapScale nodes running RAINcloudOS version 4.0.

Offering user-selectable levels of data redundancy, SnapScale uses File-level Striping to write data across multiple nodes and drives simultaneously for instant protection and high availability. With a SnapScale cluster, volumes can be configured, created, provisioned, and grown on demand. Special features such as Data Balancer redistributes files to optimize performance and Spare Distributor evenly distributes spare drives across nodes. Files can be accessed either through NFS or CIFS/SMB protocols. SnapScale Flexible Volumes automatically adjust capacity so they only occupy as much space as their data requires.



Topics in Overview:

- [SnapScale Conventions](#)
- [SnapScale Node Requirements](#)
- [RAINcloudOS Specifications](#)
- [RAINcloudOS 4.0 Features](#)
- [Client and Storage Networks](#)
- [Node Port Configurations](#)

SnapScale Conventions

The SnapScale cluster supports three or more nodes hosting redundant sets of data for data protection. An Administrator can configure, add, or remove nodes on demand to change storage requirements. The overall storage system is able to easily grow from three nodes to meet your needs.

Peer sets are created using two or three drives (based on redundancy choices) located on different nodes. Each peer set member has the same data and metadata as its peers.

There are three different states for SnapScale nodes:

- **Uninitialized node** – an independent node that has not yet been joined to a SnapScale cluster.
- **SnapScale node** – a healthy node that is a member of a fully-configured SnapScale cluster. Both 2U nodes with up to 12 drives and 4U nodes with 36 drives are available.
- **Management node** – a SnapScale node with special duties involved in managing the cluster. The Management node is selected automatically by the RAINcloudOS when the cluster boots. Should that management node fail, another currently available node is automatically chosen to become the new Management node. This Management node also hosts peer sets with metadata and data just like all other SnapScale nodes.

Other key concepts include:

- **Management IP** – the IP address through which the administrator accesses the Web Management Interface of the current Management node.
- **Peer set** – a set of two or three disks (each on a separate node) that have mirrored data for redundancy.
- **Cluster Name** – the name visible to network clients and used to connect to the cluster (similar to a server name), and resolvable to node IP addresses via round robin DNS.
- **Cluster Management Name** – the hostname resolvable to the Management IP for Web Management Interface access or Snap EDR configuration.
- **Data Replication Count** – an administrator-specified, cluster-wide count of the number of mirrored copies of data within the cluster. The Data Replication Count can be either “2” or “3” and determines the number of drives in a peer set.

A SnapScale cluster consists of two separate networks:

- **Client Network** – used exclusively for client access. Clients can connect to any node to access data anywhere on the cluster.
- **Storage Network** – an isolated network used exclusively by the cluster for inter-node communications. This includes:
 - Heartbeat (node health/presence) sensing.
 - Synchronization of peer set members.
 - Data transfer between nodes to facilitate clients reading from and writing to files.

SnapScale Node Requirements

The following table details the basic requirements for cluster nodes:

Requirement	Detailed Description
Minimum number of nodes	A SnapScale cluster must have a minimum of three (3) nodes to operate normally.
No expansion units	A SnapScale node cannot have any expansion units attached to it.
Minimum number of disks per node	Each node must have a minimum of four disks. Additional disks can be added as needed.
Maximum size of file on cluster	While the system reports total free space across the entire cluster, the maximum file size at any given time is dictated by free space on the least-utilized peer set. This is reported in the Web Management Interface.

Requirement	Detailed Description
Common Storage network	To form or join a SnapScale cluster, each Uninitialized node must be connected to the same Storage network as the other nodes.
Storage network links	To form or join a SnapScale cluster each Uninitialized node must have connectivity (active link) on both Storage network ports.
Storage network usage	Only a single cluster can use a given Storage network.
Client network separate from Storage network	The Client and Storage networks must be on different (independent) networks, and the Storage network must be isolated from all other networks.
Nodes must be running same RAINcloudOS version	To form a SnapScale cluster, all nodes must be running the same version of RAINcloudOS. To join an already configured SnapScale cluster, an Uninitialized node must have the same version of RAINcloudOS as the other SnapScale nodes: <ul style="list-style-type: none"> • If the Uninitialized node has an older version of the RAINcloudOS, the Uninitialized node must be upgraded to the later version. • If the Uninitialized node has a newer version of the RAINcloudOS, then all SnapScale nodes must be upgraded to the later version. (The node can be reinstalled with a version matching the cluster if the hardware supports it.)
Adding nodes	When adding nodes to an existing cluster, the number of nodes added at one time should be at least the same number as the Data Replication Count. This ensures the new nodes and cluster are efficiently utilizing increased storage space.
Disk requirements	All disks in the cluster must be the same type of disk (such as SAS) and same rotational speed.

RAINcloudOS Specifications

These specifications apply to all SnapScale nodes running RAINcloudOS 4.0:

Feature	Specification
Network Transport Protocols	TCP/IP (Transmission Control Protocol/Internet Protocol) UDP/IP (User Datagram Protocol/Internet Protocol)
Network File Protocols	Microsoft Networking (CIFS/SMB) UNIX Network Filesystem (NFS) 3.0 only Hypertext Transfer Protocol (HTTP/HTTPS)
Network Security	<ul style="list-style-type: none"> • Microsoft Active Directory Service (ADS) (member server) • UNIX Network Information Service (NIS) • File and Folder Access Control List (ACL) Security for Users and Groups • Secure Sockets Layer (SSL v2/3) 128-bit Encryption • SMTP Authentication and support for email encryption (STARTTLS and TLS/SSL encryption protocols)

Feature	Specification
Network Client Types	Microsoft Windows 2000 SP4/2003/2003 R2/2008 SP2/2008 R2 /XP SP3/Vista SP2/7/8/2012 Mac OS X 10.5/10.6/10.7/10.8 (via CIFS/SMB) Sun Solaris 10 and 11 HP-UX 11 AIX 5.3/6 Red Hat Enterprise Linux (RHEL) 4.x/5.x/6.x Novell SuSE Linux Enterprise Server (SLES) 10.x/11.x
Data Protection	<ul style="list-style-type: none"> • Snapshots for immediate or scheduled point-in-time images of the cluster filesystem • Support for network backup via CIFS/SMB • Support for Symantec Backup Exec 2010/2012 and NetBackup 7.5. • APC® brand Uninterruptible Power Supply (UPS) with Network Management Cards, a USB interface, or a serial interface (with USB-to-Serial adapter) are supported for graceful system shutdown
System Management	<ul style="list-style-type: none"> • Browser-based administration tool called the Web Management Interface • Read-only CLI support • Environmental monitoring • Email event notification • Data importation (migration) • SNMP (MIB II and Host Resource MIB) • User disk quotas for Windows, UNIX/Linux, FTP/FTPS • NIS Group disk quotas for UNIX/Linux
DHCP Support	Only supports Dynamic Host Configuration Protocol (DHCP) in an Uninitialized node for configuring or adding to a cluster.

RAINcloudOS 4.0 Features

NOTE: For details and descriptions of all the new features and a list of other improvements to the operating system, see the [Product Release Notes on the Overland SnapScale website](#).

With the release of RAINcloudOS 4.0, the following features and functionality are now available:

Feature	New Functionality
iSCSI Support	SnapScale can now create and host iSCSI disk targets on the cluster file system. These iSCSI disks can register with an iSNS server, and can also be managed by Windows VSS/VDS.
SMB2, FTP/FTPS, and SNMP Support Added	SMB2, FTP/FTPS, and SNMP are all now supported in RAINcloudOS.
Improved Network Monitoring	The Network Monitor page provides additional information including high-water marks, network activity for the whole cluster, and clearer labels.
Added User/Group Quotas per Volume	Storage consumption and file count quotas can now be configured for users and NIS groups per volume.

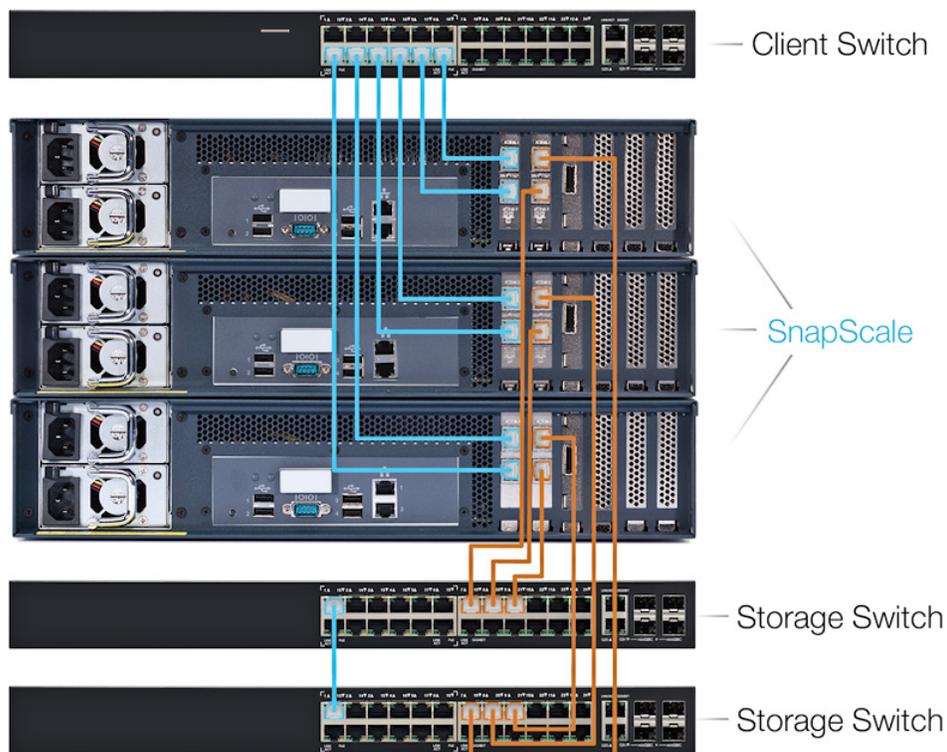
Feature	New Functionality
Data Balancer & Spare Distributor Improved	Data Balancer (formerly Capacity Balancer) redistributes files to optimize performance. Spare Distributor (formerly the Spare Disk Balancer) evenly distributes spare drives across nodes. Both have been improved for faster results.

Client and Storage Networks

SnapScale requires two separate networks to function correctly: A public network (Client) and a private network (Storage). To support failover, two Storage network switches must be connected together (using a 1GbE or 10GbE cable between the switches). Each of the two Storage network ports on the node need to be connected to a different Storage switch.

Node/Switch Cabling Example

The following example shows three dual 10GbE card X2 nodes and how to connect them to the network switches. The cables used to connect to the Client side of the network (blue) originate from the Client 10GbE card in slot 1. Two cables are used to connect both ports of each node to the Client switch.



The cables used to connect to the Storage side of the network (orange) originate from the Storage 10GbE card in slot 2. For each node, one cable is used to connect a one Storage port of each note to one of the two Storage switches used for failover.

For connections between 10GbE cards and 10GbE switches, use either direct-attached copper cables or fibre cables with SFP+ modules pre-installed in the card and switch ports.

IMPORTANT: If using fibre cables, you must use Overland-approved SFP+ modules. With the cluster powered OFF, insert the modules into the card and switch ports. Connect the fibre cable between the two SFP+ modules and restore power to the cluster.

Node Port Configurations

Both the X2 and X4 nodes come in three different configurations: 1GbE ports (both Client and Storage ports), a single 10GbE card (with 1GbE Client ports), and dual 10GbE cards.

NOTE: If desired, optional 10GbE cards can be added later to upgrade the node.

X2 Node Configurations

Basic 1GbE. At the rear of the X2 node, the 1GbE ports connected directly to the motherboard are configured to access the Client and Storage networks.



Configuration	Node GbE Ports	Network Switch
Basic 1GbE X2	Ports 1 & 2	Client (public)
	Slot 2 1GbE Card (ports 3 & 4)	Storage (private)

Single 10GbE. The single-card 10GbE configuration uses the two 1GbE ports for the Client connection and the two 10GbE ports on the card for the Storage connections.



Configuration	Node GbE Ports	Network Switch
Single 10GbE X2	Ports 1 & 2	Client (public)
	Slot 2 10GbE Card (ports 3 & 4)	Storage (private)

Dual 10GbE. The dual-card configuration uses the left 10GbE card ports for the Client connections and the right 10GbE card ports for the Storage connections. The 1GbE ports are not used.



Configuration	Node GbE Ports	Network Switch
Dual 10GbE X2	Slot 1 10GbE Card (ports 1 & 2)	Client (public)
	Slot 2 10GbE Card (ports 3 & 4)	Storage (private)

X4 Node Configurations

Basic 1GbE. At the rear of the 1GbE configuration, there are two sets of 1GbE ports connected directly to the motherboard for connecting to the switches. The top two are for the Client network; bottom two for Storage.



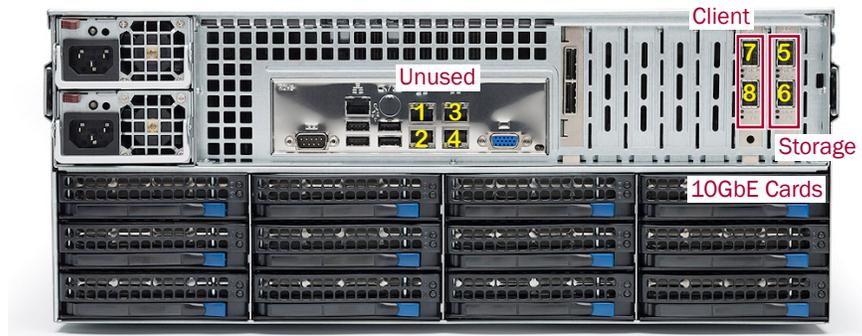
Configuration	Node GbE Ports	Network Switch
Basic 1GbE X4	Ports 1 & 2	Client (public)
	Ports 3 & 4	Storage (private)

Single 10GbE. The single-card 10GbE configuration uses the two 1GbE ports for the Client connection and the two 10GbE ports on the card for the Storage connections.



Configuration	Node GbE Ports	Network Switch
Single 10GbE X4	Ports 1, 2, 3, & 4	Client (public)
	Slot 7 10GbE Card (ports 5 & 6)	Storage (private)

Dual 10GbE. The dual-card configuration uses the left 10GbE card ports for the Client connections and the right 10GbE card ports for the Storage connections. The 1GbE ports are not used.



Configuration	Node GbE Ports	Network Switch
Dual 10GbE X4	Slot 6 10GbE Card (ports 7 & 8)	Client (public)
	Slot 7 10GbE Card (ports 5 & 6)	Storage (private)

Setup and Configuration

This section covers the initial setup and configuration of an individual SnapScale node running RAINcloudOS 4.0. It also addresses how to use that node to set up a SnapScale cluster of three or more nodes, or to add the node to an existing SnapScale cluster.

NOTE: For information concerning the installation and wiring of the SnapScale node hardware, refer to either the *SnapScale X2 Node Quick Start Guide* or the *SnapScale X4 Node Quick Start Guide*.

Topics in Setup and Configuration:

- [Connecting for the First Time](#)
- [Create a New SnapScale Cluster \(via Wizard\)](#)
- [Join an Existing SnapScale Cluster \(via Wizard\)](#)
- [Web Management Interface](#)

Connecting for the First Time

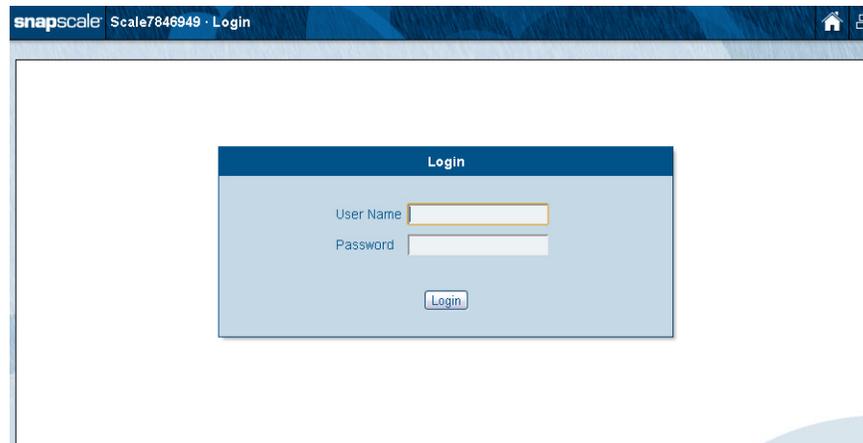
NOTE: Uninitialized nodes are configured to acquire their IP address from a DHCP server. If no DHCP server is found on the network, the node defaults to an IP address in the range of 169.254.xxx.xxx and is labeled in SnapServer Manager (SSM) as “ZeroConf”. You may not be able to see Uninitialized nodes on your network until you discover them using either the default node name or the SSM utility and optionally assign them an IP address.

Connect Using the Node Name

This procedure requires that name resolution services (via DNS or an equivalent service) be operational.

NOTE: Any node that is selected to be part of a cluster can be used to create the cluster.

1. Find the **node name** of an Uninitialized node that is to be used to create a new SnapScale cluster.
A SnapScale node name is of the format “Nodennnnnnnn,” where *nnnnnnnn* is the node chassis number. The node number is a unique, numeric-only string that appears on a label affixed to the bottom of the appliance.
2. In a web browser, enter the **URL** to connect to the node.
For example, enter “http://Nodennnnnnnn” (using the node name).
3. Press **Enter** to connect to the Web Management Interface.



4. In the login dialog box, enter **admin** as the user name and **admin** as the password (the system defaults), then click **OK**.
5. Complete the **Initial Setup Wizard** to either create a new SnapScale cluster or join an existing cluster.

Connect Using SSM

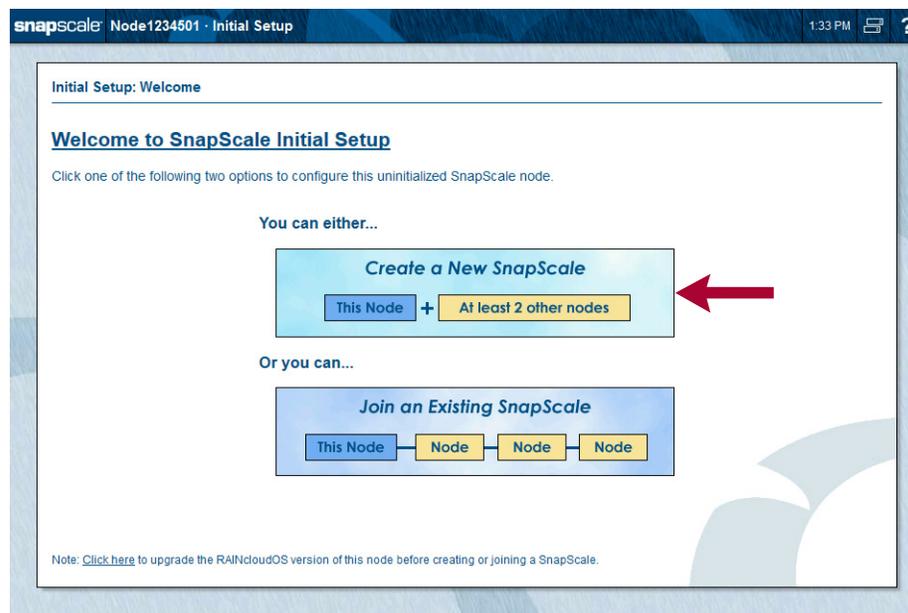
1. Launch **SnapServer Manager (SSM)**.
SSM discovers all SnapServers, SnapScale clusters, and SnapScale nodes on its local network segment and displays their names, IP addresses, and other information in the main console. If you do not have a DHCP server, there might be a delay before the node appears on the network.

NOTE: To distinguish multiple SnapServers or SnapScale nodes, you may need to find their default names as explained in [Connect Using the Node Name](#) on page 2-1.
2. If using a DHCP server, proceed to [Step 3](#); otherwise, assign an **IP address** to one of the nodes to be configured in the cluster.

NOTE: Only one node needs to be configured with an IP address in order to create the cluster.
 - a. In SSM, right-click the **node name**.
 - b. Select **Set IP Address**.
 - c. Enter an IP address and a subnet mask, then click **OK**.
3. In SSM, right-click the node name and select **Launch Web Administration**.
4. Log into the Web Management Interface.
In the login dialog box, enter **admin** as the user name and **admin** as the password (the system defaults), then click **OK**.
5. Complete the **Initial Setup Wizard** to either create a new SnapScale cluster or join an existing cluster.

Create a New SnapScale Cluster (via Wizard)

On a new node, once you log in, the Initial Setup Wizard runs displaying the **Welcome** page. From the Initial Setup Wizard, you can use this node to create a new SnapScale cluster by connecting to two or more other nodes. Click **Create a New SnapScale** to start the wizard.



The Initial Setup Wizard for **creating** a new SnapScale cluster consists of seven steps:

Step 1: Select the nodes to be included in the cluster.

Step 2: Review the Client network information.

Step 3: Choose the static TCP/IP settings for the Client network.

Step 4: Populate the Static IP addresses for the nodes.

Step 5: Enter the basic SnapScale properties.

Step 6: Set the date and time.

Step 7: Verify the settings and create a SnapScale cluster.

NOTE: After the cluster is created, you are asked to configure the Administrator's password as part of Step 7.

Step 1 – Select SnapScale Nodes

Select the nodes you want to use from the list of eligible nodes.

 **IMPORTANT:** At least three nodes are required to create a SnapScale clustered network. All nodes must have the identical version of RAINcloudOS (ROS) and be on a subnet that does not contain an existing cluster. The Client network interfaces for all the nodes must be located on the same public network subnet, and the Storage network interfaces for all nodes must be located on the same private Storage network subnet. The nodes cannot have any expansion units attached.

Any combination of node types (X4 and X2) can be used to create a cluster.

Initial Setup: Create SnapScale - Select SnapScale Nodes

Select the nodes below that you want to add to this new SnapScale and click Next. (All eligible nodes are selected by default.)

Note: 3 nodes are required as a minimum to create a SnapScale. All nodes in a SnapScale must have identical RAINcloudOS (ROS) versions, and the client network interface for all nodes must be located on the same subnet.

13 Eligible Nodes. (This node: ROS version=4.0.119, IP address=192.168.199.101, Subnet mask=255.255.254.0)

Node	Model	ROS Version	Disks	Add to SnapScale
Node1234501 (This Node)	X2	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
Node1234503	X2	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
Node1234507	X2	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
Node1234510	X4	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk) 13: 1.95 TB 14: 1.95 TB 15: 3.91 TB 16: 3.91 TB 17: 1.95 TB 18: 1.95 TB 19: 3.91 TB 20: 3.91 TB 21: 1.95 TB 22: 1.95 TB 23: 3.91 TB 24: 3.91 TB Rear Disks: 25: 1.95 TB 26: 1.95 TB 27: 3.91 TB 28: (No Disk) 29: (No Disk) 30: (No Disk) 31: (No Disk) 32: (No Disk)	<input checked="" type="checkbox"/>

Back Re-Detect Available Nodes Next

Verify that the boxes in the Add to SnapScale column for the nodes you want to use are checked. Click **Re-Detect Available Nodes** to refresh the list. When ready, click **Next**.

NOTE: If you deselect one or more of the detected nodes, when you click Next a message page is displayed recommending that you add all the nodes at once.

Initial Setup: Create SnapScale - Select SnapScale Nodes

Important: You have selected 3 out of 6 eligible nodes for your new SnapScale. Overland recommends that you select all eligible nodes now for addition to the SnapScale, rather than adding them after the SnapScale is created.

Click Keep My Node Selections if you wish to ignore the above concern(s) and continue Initial Setup using the nodes you have selected.

Back Keep My Node Selections

Step 2 – Client Network Configuration Overview

Review the information about setting up your Client network. Click **Next** to continue.

Initial Setup: Create SnapScale - Configure Client Network : Overview

Step 1 Step 2 Step 3 Step 4 Step 5 Step 6 Step 7

Your SnapScale nodes are connected together via a *storage network* that allows both user data and meta-data to be communicated between nodes. This storage network is automatically configured for you as part of this Initial Setup process. Your SnapScale also includes a *client network* which will be used by users to access (read & write) user data stored on the SnapScale. This client network is configured and managed by you using static IP addresses.

SnapScale Network Overview

Client Network (Public)

Node Node Node Node

Storage Network (Private)

Back Next

(Click Next to configure the client network settings for your SnapScale.)

Step 3 – Choose Client Network Static TCP/IP Settings

Use this step to specify the static TCP/IP settings that will be common to all nodes in the cluster. Then click **Next** to continue to the next page to set the actual node static IP addresses.

Initial Setup: Create SnapScale - Configure Client Network : Static TCP/IP Settings

Step 1 Step 2 Step 3 Step 4 Step 5 Step 6 Step 7

Specify the static TCP/IP settings that are common to all nodes in the SnapScale. In the next step you will specify a list of static IP addresses to be used by the nodes.

Client network static TCP/IP settings.

Subnet Mask: 255.255.0.0

WINS Servers: (optional)

Default Gateway: 10.25.1.1 (optional)

DNS Domain Name: devnet.myoverland.net (optional)

Domain Name Servers: 10.6.8.34 (optional), 10.6.8.35 (optional)

Back Next

Step 4 – Configure Node Static IP Addresses

A SnapScale cluster requires a set of static IP addresses: one for each node, and one for the Management IP. Use this page to specify the static IP addresses for each of your nodes and for the SnapScale Management IP address used to access the Web Management Interface for this cluster.

snap scale Node2413872 · Initial Setup 4:19 PM ?

Initial Setup: Create SnapScale - Configure Client Network : Static Node IP Addresses

Step 1 Step 2 Step 3 **Step 4** Step 5 Step 6 Step 7

Specify 4 static IP addresses: one for each of your SnapScale nodes and one for the SnapScale Management IP address (the Management IP address is used to access the Web Management Interface for this SnapScale). These IP addresses must all be located on the same subnet, and they will be automatically assigned to your nodes when the SnapScale is created.

Optional: Enter a starting IP address and click the "Populate" button to populate the list below with sequential static IP addresses. You can then review or change the IP addresses before clicking Next.

Starting IP Address (optional)

Enter static IP addresses below:

Static IP Address	
<input type="checkbox"/>	<input type="text"/> (Management IP address.)
<input type="checkbox"/>	<input type="text"/> (Node IP address.)
<input type="checkbox"/>	<input type="text"/> (Node IP address.)
<input type="checkbox"/>	<input type="text"/> (Node IP address.)

These IP addresses must all be located on the same subnet. They are automatically assigned to your nodes when the SnapScale cluster is created.

The **Populate Static IP Addresses** button can be used to automatically enter a sequential list of static IP addresses. Just enter an IP address on the subnet and click **Populate Static IP Addresses**. The fields below it are automatically populated.

snap scale Node2413872 · Initial Setup 4:21 PM ?

Initial Setup: Create SnapScale - Configure Client Network : Static Node IP Addresses

Step 1 Step 2 Step 3 **Step 4** Step 5 Step 6 Step 7

Specify 4 static IP addresses: one for each of your SnapScale nodes and one for the SnapScale Management IP address (the Management IP address is used to access the Web Management Interface for this SnapScale). These IP addresses must all be located on the same subnet, and they will be automatically assigned to your nodes when the SnapScale is created.

Optional: Enter a starting IP address and click the "Populate" button to populate the list below with sequential static IP addresses. You can then review or change the IP addresses before clicking Next.

Starting IP Address (optional)

Enter static IP addresses below:

Static IP Address	
<input type="checkbox"/>	10.25.11.100 (Management IP address.)
<input type="checkbox"/>	10.25.11.101 (Node IP address.)
<input type="checkbox"/>	10.25.11.102 (Node IP address.)
<input type="checkbox"/>	10.25.11.103 (Node IP address.)

Click **Next** to continue.

Step 5 – Basic SnapScale Properties

Use this step to enter the basic properties for your new SnapScale cluster, then click **Next**.

Initial Setup: Create SnapScale - Basic SnapScale Properties

Step 1 Step 2 Step 3 Step 4 **Step 5** Step 6 Step 7

Enter basic properties for the new SnapScale and click Next.

SnapScale name and description.
 SnapScale Name
 SnapScale Description (optional)

Data replication count.
 The data replication count specifies how many copies of each data file or folder to maintain. A replication count of 3x offers higher data protection yet uses more disk space.
 Data Replication Count (Note: Once the SnapScale is created, the replication count can be changed only from 3x to 2x.)

Spare disks.
 Spare disks specifies the number of available disks in the SnapScale to reserve for spares. A spare disk is used to automatically replace a failed peer set member.
 Allocate spare disks
 Spare Disks

Snapshots.
 If you plan on using snapshots, it is recommended that you reserve at least 20% of your SnapScale storage space for snapshots.
Note: Once the SnapScale is created, the storage space reserved for snapshots can be decreased, *but never increased*.
 Reserve space for snapshots
 Percentage of SnapScale storage to reserve for snapshots

This table lists and describes the basic options:

Option	Description
SnapScale Name	<p>Either accept the default name or enter an alphanumeric name up to 15 characters in length. Network clients use this name with round robin DNS name resolution to connect to the cluster.</p> <p>The default name is "Scalennnnnn" (where nnnnnn is the appliance number of the node used to create the cluster).</p>
SnapScale Description	<p>This optional field provides a place to define the cluster in the overall scheme of your network and better identify the cluster on a LAN.</p>
Data Replication Count	<p>The data replication count establishes the level of data redundancy in the cluster. The setting specifies how many disks are in a peer set and as a result how many copies of each data file or folder to maintain. A count of 3x offers higher data protection but uses more disk space.</p> <p>Once the cluster is created, the count can only be decreased from 3x to 2x. It cannot be increased from 2x to 3x.</p>
Spare Disks Allocation	<p>Check the box and select the number of spare disks you want to reserve. A spare disk is used to automatically replace a failed Peer Set member.</p> <p>If there are unused drives remaining after allocating the number of spares requested, they are used for other peer sets. If there is an insufficient number of drives left to create a final peer set, the drives are configured as additional spares.</p>

Option	Description
Reserve Space for Snapshots	<p>Check the box and select the percentage of the storage space you want to reserve for snapshots. It is recommended that at least 20% of your SnapScale storage space be set aside for snapshots.</p> <p>NOTE: Once the SnapScale cluster is created, the storage space reserved for snapshots can only be decreased. It can never be increased.</p>



IMPORTANT: If you uncheck the box for reserving space for snapshots, an alert is displayed to remind you that the feature will be permanently disabled for the cluster.



Step 6 – Set Date and Time

Nodes automatically synchronize time with one another. You can either manually set the date and time to specific values, or you can use NTP (Network Time Protocol) servers to automatically synchronize the date and time. Visit www.ntp.org for a list of public NTP primary and secondary servers, or simply use the default NTP servers below.

If you intend to join the cluster to a Windows domain, configure the cluster using the manual settings to set the date and time. Otherwise, configure the cluster to synchronize with up to two NTP servers.

NOTE: NTP cannot be used if you are joining a Windows Active Directory domain.

Default NTP servers automatically populate the server fields. The Time Zone is set automatically to UTC time but can be changed using the drop-down list.

Click **Next** to continue.

Step 7 – Summary Verification & Cluster Creation

At this step, review the current settings and go back if you need to make changes.

NOTE: Make note of the Management IP address for later use. Also, both the Client and Storage network bond types can be changed after the cluster is created. See TCP/IP Networking in Chapter 4.

Initial Setup: Create SnapScale - Create SnapScale Summary

Step 1 | Step 2 | Step 3 | Step 4 | Step 5 | Step 6 | **Step 7**

Please review your settings below and click Create New SnapScale to complete the creation of this SnapScale.

Important Security Note: You will be asked to change your administrator password after the SnapScale has been successfully created. The administrator password is used to access this Web Management Interface.

SnapScale settings.

SnapScale Name	MyCluster
Data Replication Count	2x
Spare Disks	2
Snapshot Space Reserved	20%
Management IP Address	192.168.199.100 (Please make note of this IP address for later use.)
Subnet Mask	255.255.0.0
Default Gateway	192.168.192.1
Domain Name Servers	192.168.192.22
WINS Servers	-
DNS Domain Name	devnet.myoverland.net
Time Zone	(UTC-08:00) Pacific Time (US & Canada)

3 SnapScale Nodes. (Note: The IP addresses displayed below will not necessarily be assigned to their associated node.)

Node	IP Address	Model	ROS Version	Disks											
Node1234501 (This Node)	192.168.199.101	X2	4.0.119	1: 1.95 TB	2: 1.95 TB	3: 1.95 TB	4: 3.91 TB	5: 1.95 TB	6: 1.95 TB	7: 3.91 TB	8: 3.91 TB	9: 1.95 TB	10: (No Disk)	11: (No Disk)	12: (No Disk)
Node1234503	192.168.199.102	X2	4.0.119	1: 1.95 TB	2: 1.95 TB	3: 1.95 TB	4: 3.91 TB	5: 1.95 TB	6: 1.95 TB	7: 3.91 TB	8: 3.91 TB	9: 1.95 TB	10: (No Disk)	11: (No Disk)	12: (No Disk)
Node1234507	192.168.199.103	X2	4.0.119	1: 1.95 TB	2: 1.95 TB	3: 1.95 TB	4: 3.91 TB	5: 1.95 TB	6: 1.95 TB	7: 3.91 TB	8: 3.91 TB	9: 1.95 TB	10: (No Disk)	11: (No Disk)	12: (No Disk)

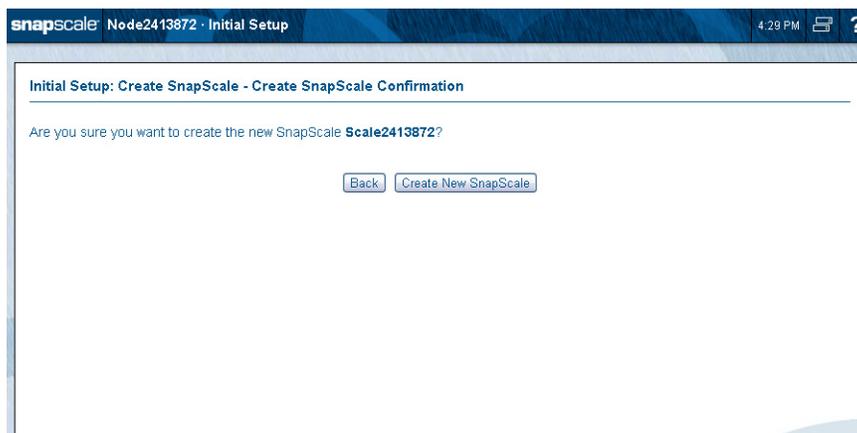
Buttons: Back | Create New SnapScale

IMPORTANT: If you uncheck the box for reserving space for snapshots, an alert is displayed to remind you that the feature will be permanently disabled for the cluster.

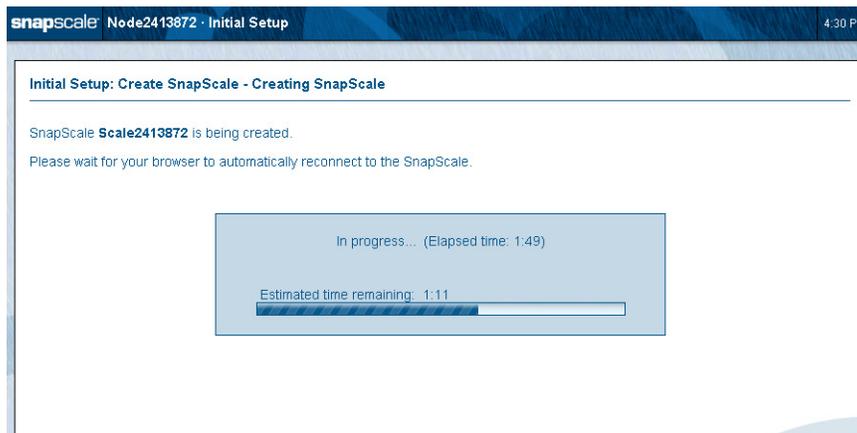
SnapScale settings.

SnapScale Name	SimScale
Data Replication Count	2x
Spare Disks	2
Snapshot Space Reserved	Warning: No space is being reserved for snapshots. The Snapshots feature will be permanently disabled for this SnapScale.
Management IP Address	192.168.199.100 (Please make note of this IP address for later use.)
Subnet Mask	255.255.254.0

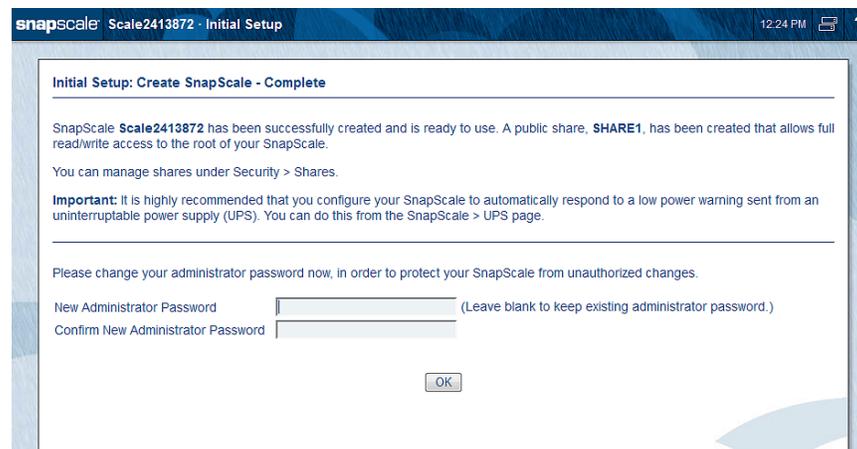
Click **Create New SnapScale** to complete the process. A confirmation page is shown.



Click **Create New SnapScale** again to create the cluster. A progress bar is displayed as the SnapScale cluster is created.

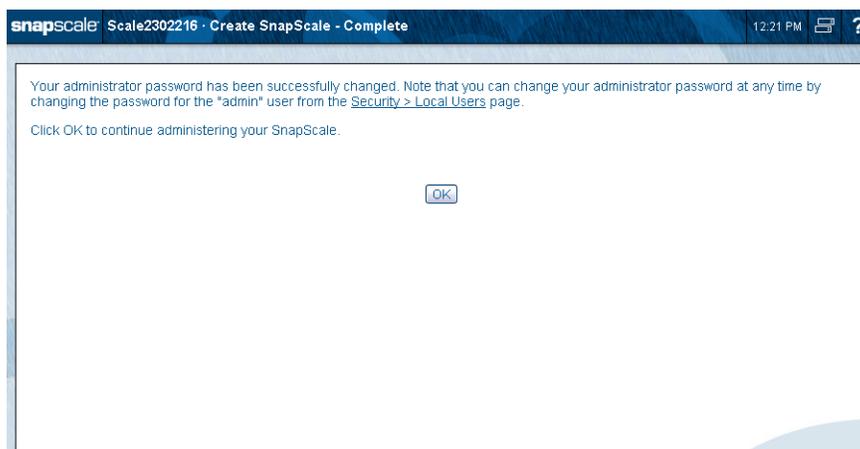


Once the cluster is created and the system changes the uninitialized node IP addresses from DHCP to the configured static IP address, a completion page is displayed stating that a share was created and suggesting UPS units be enabled. To enhance security, you are asked to change the default administrator password after the cluster has been successfully created:



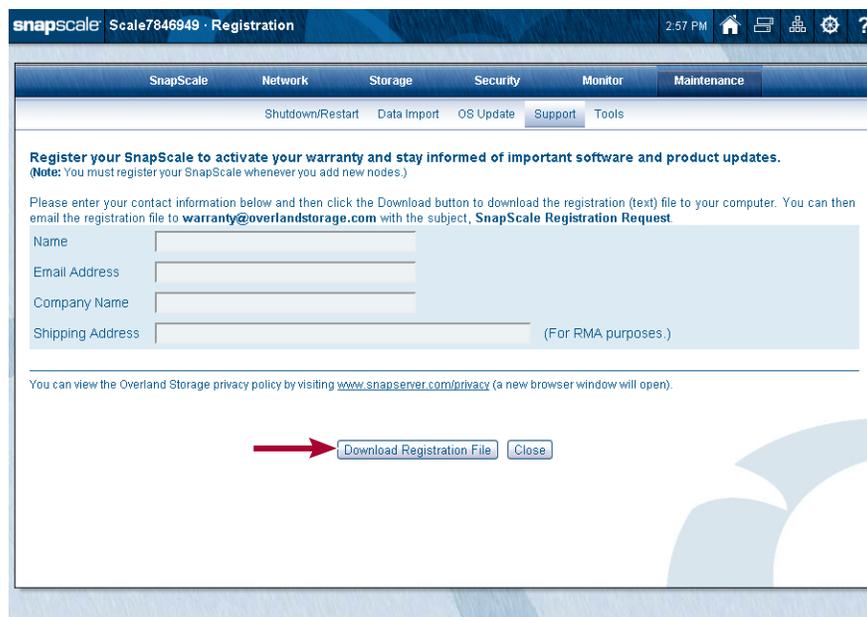
It is highly recommended that you use the password fields at the bottom of the page to change the Administrator's password for the cluster.

After changing the Administrator's password and clicking **OK**, a success page is shown:



Click **OK** to continue. The **Login** page is shown. Log in using the new password.

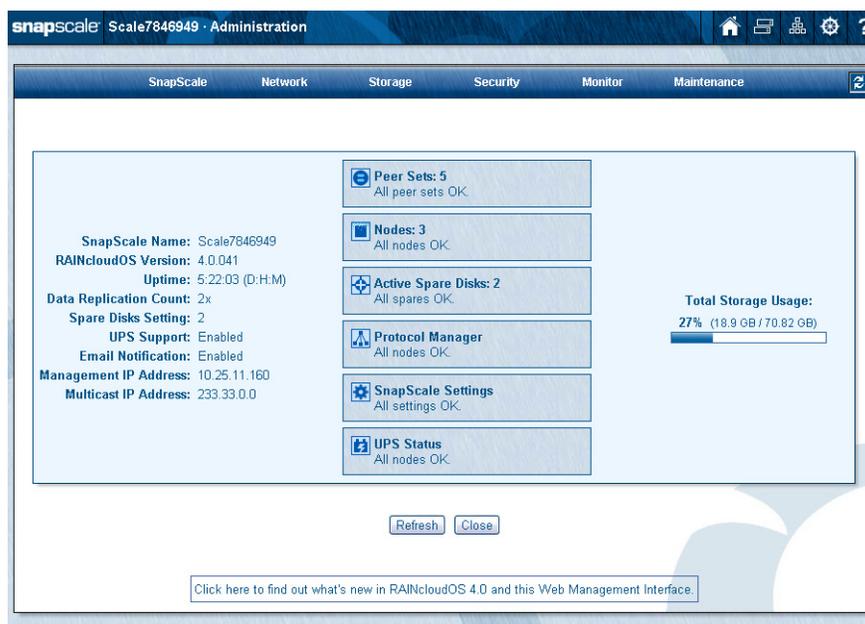
After changing the password and logging back in, the **Registration** page is displayed to facilitate activating your warranty:



Complete the registration fields and then click **Download Registration File**. Email that file (SnapScaleRegistration.csv) to Overland Storage Service (warranty@overlandstorage.com) using the subject line "SnapScale Registration Request" to initiate your warranty coverage. (See [To Register Your Cluster](#) in Chapter 8.)

Click **Close**. You will receive a confirmation email to confirm and complete the registration.

When you close that page, the **Administration** page is displayed:



It is recommended that you configure your DNS in your network so clients can resolve the cluster using round-robin name resolution:

- Add a host record for the cluster management name (<clustername>-mgt) to resolve to the Management IP address.
- Add multiple host records for the cluster name resolving to each of the node IP addresses. The DNS resolves lookups for the cluster name via round robin.

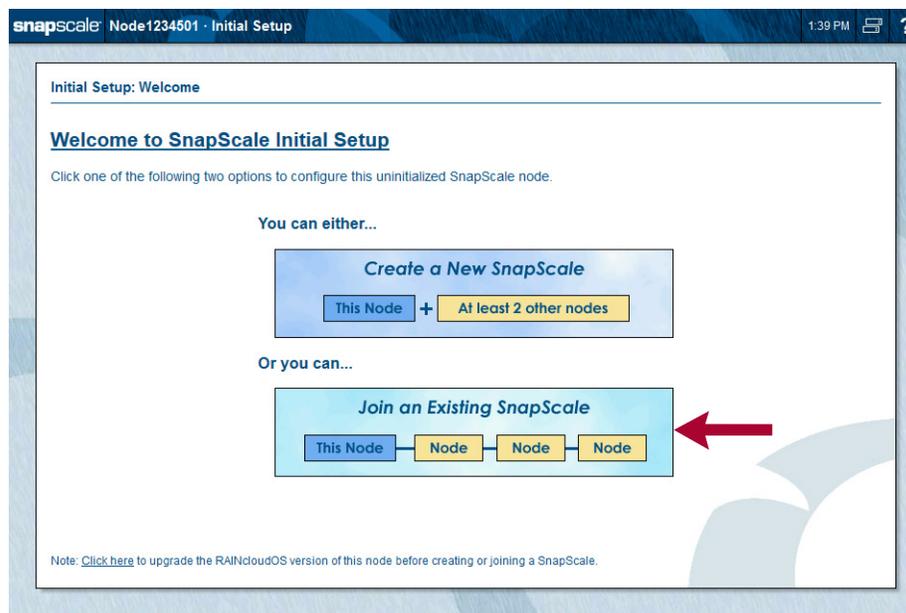
Join an Existing SnapScale Cluster (via Wizard)

 **IMPORTANT:** While the Initial Setup Wizard can be used to add one or more new nodes to an existing cluster, it is recommended that you log into the existing cluster's Web Management Interface and add the nodes using the Add Nodes function (**Storage > Nodes > Add Nodes**). Refer to [Adding Nodes](#) in [Chapter 5](#) for more information.

At any time, one or more new nodes can be added to the cluster to expand the storage pool.

NOTE: To create new peer sets to expand cluster storage, it is recommended that the number of new nodes you add is equal to the Data Replication Count being used (2x or 3x) and they all be added at the same time.

When you log into any of the new, uninitialized nodes, the Initial Setup Wizard launches displaying the **Welcome** page and its two options. To add this and other nodes to an existing SnapScale cluster, click **Join an Existing SnapScale**.



The Initial Setup Wizard then redirects you to the **Add Nodes** page in the Web Management Interface where this node (and all other discovered/new nodes) can be easily added to the cluster. (See [Adding Nodes](#) in [Chapter 5](#) for more information.) You are then directed to select the nodes to add, set the static IP addresses, and confirm the settings.

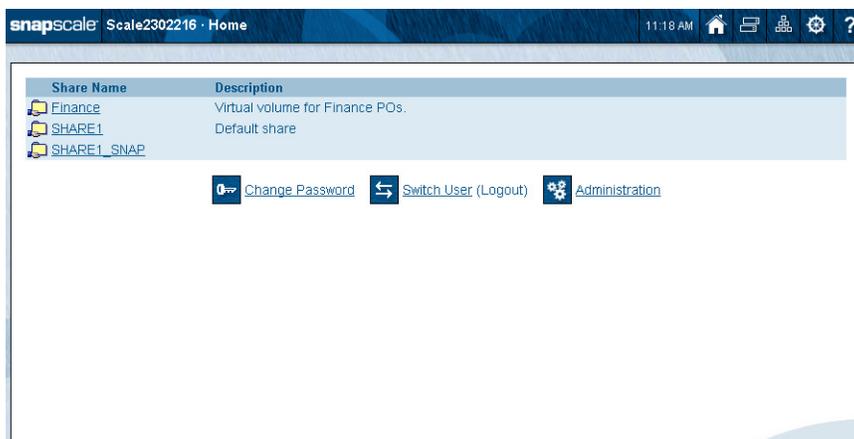
NOTE: If no existing SnapScale cluster is detected, a warning is displayed. Verify that the node is on the same Storage network as the other nodes in the cluster, then click Re-Detect SnapScale.

Web Management Interface

SnapScale nodes use a web-based graphical user interface (GUI), called the Web Management Interface, to administer and monitor the cluster. It supports most common web browsers. JavaScript must be enabled in the browser for it to work.

When connecting to the cluster with a web browser, the Web Home page (see [Web Home](#) in [Chapter 9](#)) of the Web Management Interface is displayed. This page shows any shares at the top, the three primary options below the shares list, and has special navigation buttons displayed on the right side of the title bar (see the next table).

NOTE: If you have not gone through the initial setup or authentication is required, you may be prompted to log in when you first access the Web Management Interface.



The Web **Home** page displays the following icons and options:

Icons & Options	Description
Change Password 	Click this icon to access the password change page. Passwords are case sensitive. Use up to 15 alphanumeric characters.
Switch User 	Click this icon to log out and open the login dialog box to log in as a different user.
Administration 	Click this icon to administer the node. If you are not yet logged in, you are prompted to do so.
Navigation Buttons:     	The following navigation buttons are present in the upper right on every Web Management Interface page: Home – Click this icon to switch between the Web Home page and the Admin Home page. If you have not yet logged in to the Admin Home page, only the Web Home page is available. Snap Finder – Click this icon to view a list of all SnapServers, SnapScale clusters, and Uninitialized nodes on your network, and to specify a list of remote servers that can access these servers, clusters, and nodes on other subnets. You can access these servers, clusters, and nodes by clicking the listed name or IP address. SnapExtensions – Click this to view the SnapExtensions page, where you can acquire licenses for and configure third-party applications. Site Map – Click this icon to view a Site Map of the available options in the Web Management Interface, where you can navigate directly to all the major utility pages. The current page is shown in orange text. Help – Click this icon to access the web online help for the Web Management Interface page you are viewing.

Icons & Options	Description
UI Appearance	Click the Mgmt. Interface Settings link in the Site Map to choose a background for the Web Management Interface. You can select either a solid-colored background or a textured-graphic background.

When logged in to the **Administration** page, details about the cluster's health are shown:

The screenshot displays the SnapScale Administration interface. At the top, there is a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. The main content area is divided into two columns. The left column contains system information: SnapScale Name (Scale7846949), RAINcloudOS Version (4.0.041), Uptime (5:22:03 (D:H:M)), Data Replication Count (2x), Spare Disks Setting (2), UPS Support (Enabled), Email Notification (Enabled), Management IP Address (10.25.11.160), and Multicast IP Address (233.33.0.0). The right column features several status boxes: Peer Sets (5, All peer sets OK), Nodes (3, All nodes OK), Active Spare Disks (2, All spares OK), Protocol Manager (All nodes OK), SnapScale Settings (All settings OK), and UPS Status (All nodes OK). A Total Storage Usage bar shows 27% (18.9 GB / 70.82 GB). At the bottom, there are Refresh and Close buttons, and a link to find out what's new in RAINcloudOS 4.0 and this Web Management Interface.

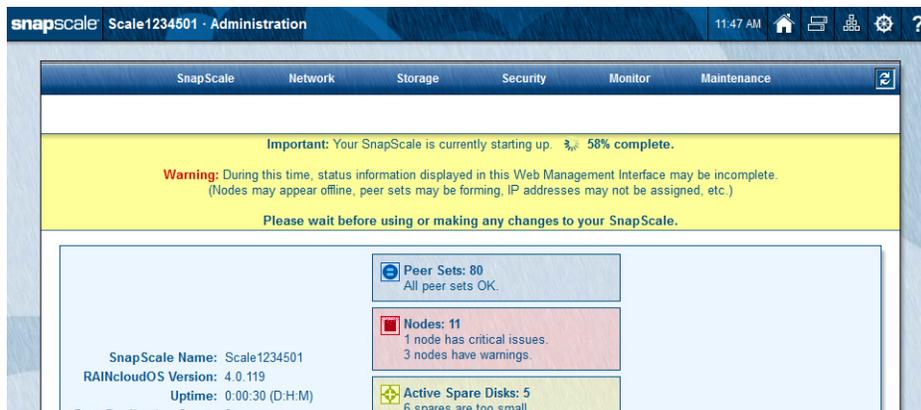
The same icons are available at the top of the page plus a refresh icon (🔄) for auto-refresh pages located on the tab bar. For more information, see [Web Home](#) in [Chapter 9](#).

Alert Messages

Alert messages are displayed on Administrator-level Web Management Interface pages that display a menu. Some alerts (such as Spare Distributor and Data Balancer) have clickable options:

- **[Later]** - Hides the alert for 24 hours or until after feature is run, whichever is first.
- **[Hide]** - Suppresses the alert. It will not be shown again until after the feature called out in the alert is run and a new alert for that feature is generated.

When a cluster is restarted, the Web Management Interface shows the status while the cluster is booting. Because some components are not immediately available, an alert message is displayed showing the percent done and as a reminder that the process is not complete, some nodes may appear offline, and so forth. Some of the status boxes may show warnings.



Site Map

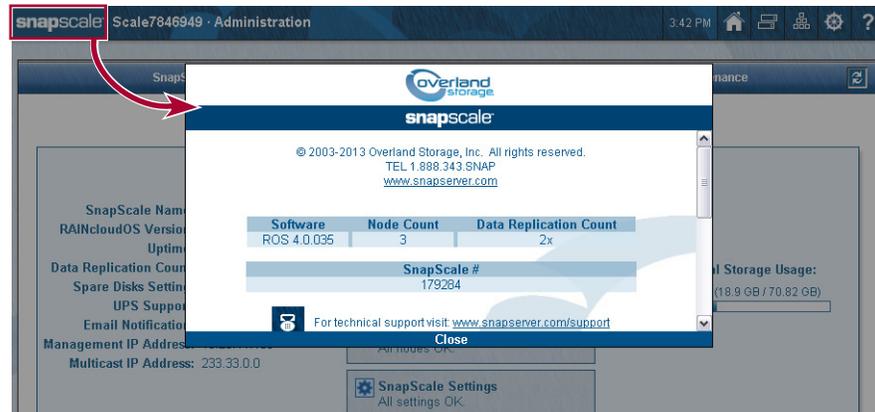
The RAINcloudOS site map (⚙️) provides links to all the web pages that make up the Web Management Interface. All the pages are each covered in detail in the following chapters.



To close the site map, click either **Close** or outside the map.

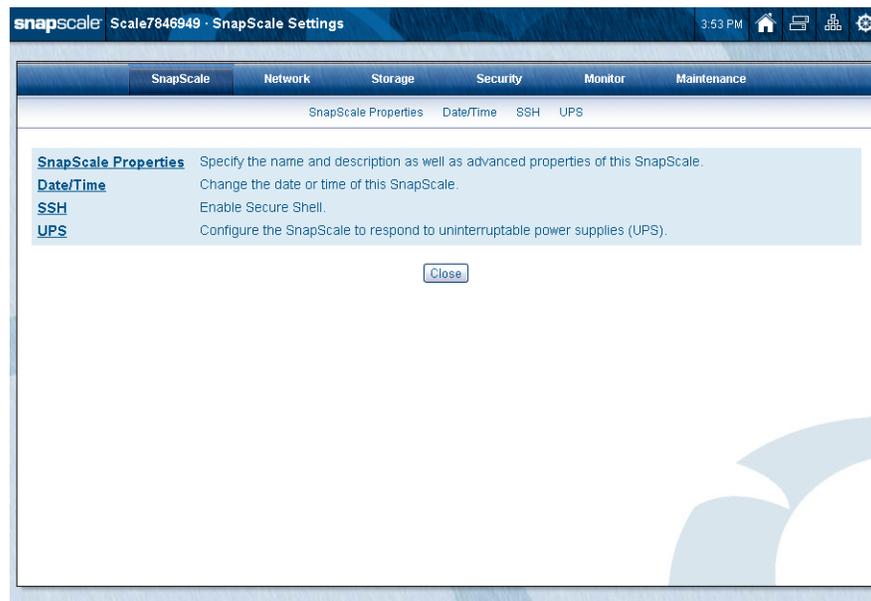
Contact Information

From the Web Management Interface, click the SnapScale logo in the upper left corner of the Web Management Interface to display the pertinent hardware, software, and contact information:



Scroll down to view additional contact information. Click either **Close** or outside the box to dismiss.

This section covers the initial setup and configuration of a SnapScale cluster of three or more nodes. The four basic options for cluster settings are found under the SnapScale tab. They can also be accessed using the site map icon (⚙️).



Topics in SnapScale Settings:

- [SnapScale Properties](#)
- [Date/Time](#)
- [SSH](#)
- [UPS](#)

SnapScale Properties

These basic options are found under **SnapScale Properties**:

The screenshot shows the SnapScale Properties configuration interface. At the top, there's a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Below this, there are sub-tabs for SnapScale Properties, Date/Time, SSH, and UPS. The main content area contains several configuration sections:

- SnapScale Name:** A text input field containing "Scale7846949".
- Description:** An optional text input field.
- Data Replication Count:** A dropdown menu set to "2x". A note below it states: "(Note: The replication count cannot be changed.)"
- Spare Disks:** A checkbox labeled "Allocate spare disks" is checked. Next to it are "Active spares" with a plus icon, "1" with a minus icon, and "2" with a plus icon. Below this is a dropdown menu set to "2".
- Storage Usage Warning Percentage:** A dropdown menu set to "80".
- Storage Usage Critical Percentage:** A dropdown menu set to "95".

At the bottom of the form are "OK" and "Cancel" buttons.

This table details the options on the **SnapScale Properties** page:

Option	Description
SnapScale Name and Description	<p>Either accept the default cluster name or enter an alphanumeric name up to 15 characters in length. Network clients can use this name along with round robin DNS name resolution to connect to the cluster.</p> <p>The default name is "Scalennnnnn" (where nnnnnn is the appliance number of the node used to create the cluster).</p>
Description	<p>This optional field provides a place to define the cluster in the overall scheme of your network and better identify the cluster on a LAN.</p>
Data Replication Count	<p>The data replication count establishes the level of data redundancy in the cluster. The setting specifies how many copies of each data file or folder to maintain. A count of 3x offers higher data protection but uses more disk space.</p> <p>Once the cluster is created, the count can only be decreased from 3x to 2x. It cannot be increased from 2x to 3x.</p>
Spare Disks	<p>Check the box and select the number of spare disks you want to reserve. A spare disk is used to automatically replace a failed Peer Set member.</p> <p>If there are unused drives remaining after allocating the number of spares requested, they are used for other peer sets. If there is an insufficient number of drives left to create a final peer set, the drives are configured as additional spares.</p>
Storage Utilization	<p>Use the two drop-down lists to select the percentage of storage used before a warning or critical notice is sent.</p> <p>If not done already, use the link in this section to set up email notification. See Email Notification in Chapter 8.</p>

Date/Time

You can set the cluster date and time manually or have it set automatically via NTP or Windows Active Directory domain membership. Nodes automatically synchronize time with one another.

An ISO 8601 time stamp is applied when recording node activity in the Event Log (**Monitor** tab), when creating or modifying files, and when scheduling snapshot operations. Use this page to configure date and time settings:

CAUTION: If the current date and time are reset to an earlier date and time, the change does not automatically propagate to any scheduled events you have already set up for snapshot or Snap EDR operations. These operations continue to run based on the previous date and time setting. To synchronize these operations with the new date and time settings, you must reschedule each operation.

Configure Date and Time Settings Manually

1. Click the **Set the date and time** button.
2. Edit date and time **settings** as described in this table:

Option	Description
Date	Enter the current date in the format indicated.
Time	Enter the current time in the format indicated.
Time Zone	Select the time zone that you want to use for this node.

3. From the drop-down list, select the **Time Zone** for the cluster.
4. Click **OK** when finished.

Once you join a Windows domain, the settings are automatically adjusted to synchronize with the domain settings.

NOTE: RAINcloudOS automatically adjusts for Daylight Saving Time, depending on your time zone.

Configure Date and Time Settings for Automatic Synchronization

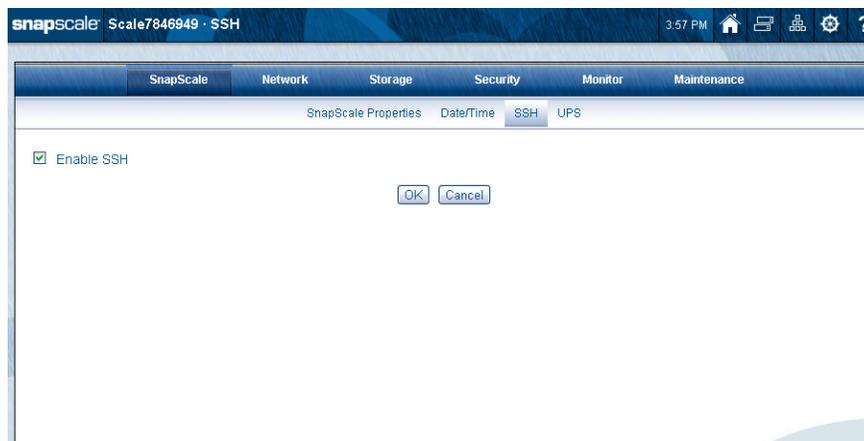
If the cluster is not joined to a Windows Active Directory domain, you can use the automatic synchronization option to configure the cluster to set date and time automatically via Network Time Protocol (NTP).

1. Click the **Automatically Synchronize** button.
 - Default NTP servers are displayed. To accept them, skip to [Step 2](#).
 - Otherwise:
 - Enter the **address** for the primary NTP server.
 - Optionally, enter a **second IP address** for a different NTP server as backup.
2. From the drop-down list, select the **Time Zone** for the cluster.
3. Click **OK** when finished.

In some cases, this change may require you to log back in to the Web Management Interface.

SSH

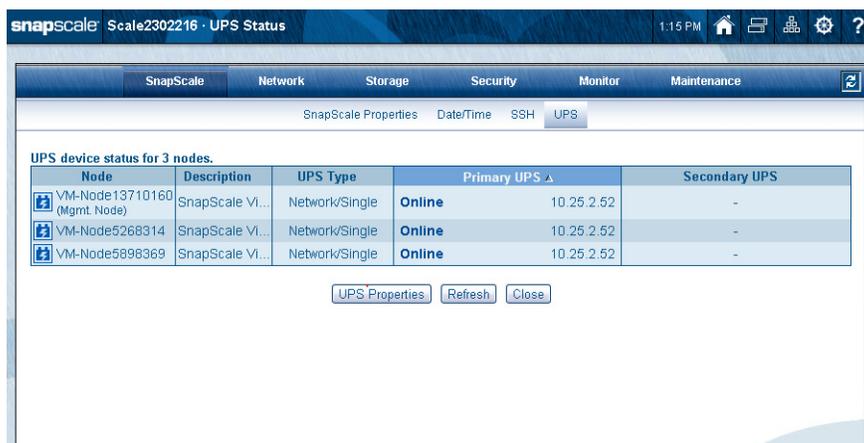
This page provides the ability to enable/disable Secure Shell (SSH) on the cluster for security purposes. By default, it is enabled.



UPS

SnapScale supports automatic shutdown when receiving a low-power warning from an APC uninterruptible power supply (UPS). Use **SnapScale > UPS** to manage this feature:

NOTE: If UPS devices have not been configured, the first time you select this option, you are automatically shown the UPS Properties page. See [Edit UPS Properties](#) on page 3-6.



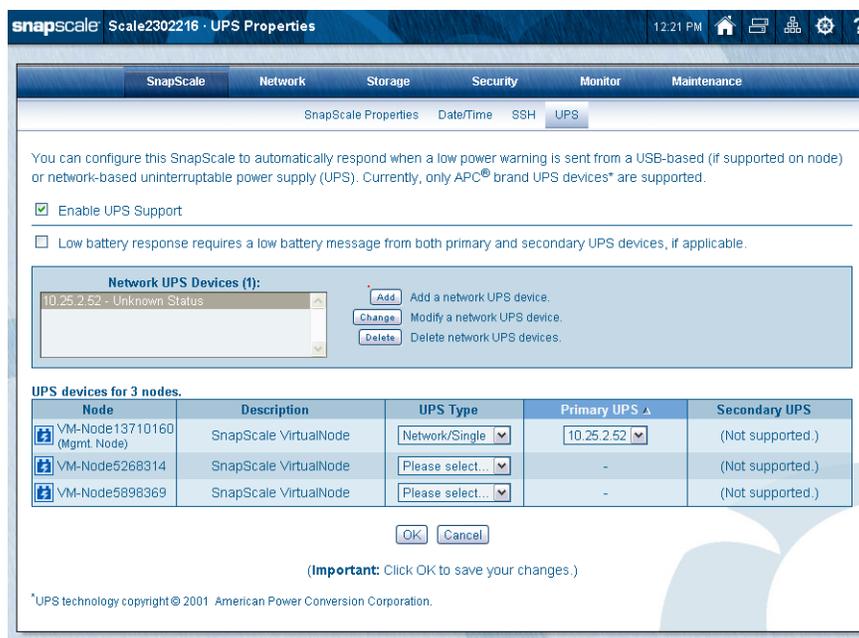
An APC Smart-UPS[®] series device allows the SnapScale cluster to shut down gracefully in the event of an unexpected power interruption. You can configure the cluster to automatically shut down when a low power warning is sent from one or more APC network-enabled or USB-based UPS devices (some serial-only APC UPS devices are also supported by using the IOGear GUC232A USB to Serial Adapter Cable). To do this, you must enable UPS support on the cluster, as described in this section, to listen to the IP address of one or more APC UPS devices, and you must supply the proper authentication phrase configured on the UPS devices.

NOTE: Select a UPS capable of providing power to a SnapScale node for at least ten minutes. In addition, in order to allow the cluster sufficient time to shut down cleanly, the UPS must be configured to provide power for at least five minutes after entering a low battery condition.

Edit UPS Properties

To manage the network UPS devices, click the **UPS Properties** button:

NOTE: If UPS devices have not been configured, the first time you select that option, you are automatically shown the UPS Properties page.



UPS Properties page options:

Option	Description
Enable UPS Support	Check the Enable UPS Support box to enable support.
Low battery response message	Check the box to initiate a graceful shutdown only when both the primary and secondary UPS devices for a node send a low battery message.
Network UPS Devices (#)	This field shows a list of UPS devices that are used with the cluster. Use the Add , Change , and Delete buttons to manage the list.
UPS Type (Third column in Node table)	Use the drop-down list in the third column of the Node table to select which UPS device is used: <ul style="list-style-type: none"> • USB – Select this option to use a direct-attached (USB) device. • Network/Single – Use this option to select a network UPS device. • Network/Dual – Use this option to activate the option of a secondary network UPS device.

Option	Description
Primary UPS (Fourth column in Node table)	Selecting the Network/Single option under UPS Type causes a drop-down list to be displayed in this column. Select the primary UPS to associate with the node from the list (which is based on the Network UPS Devices table).
Secondary UPS (Fifth column in Node table)	If supported, selecting the Network/Dual option (under UPS Type) causes a drop-down list to be displayed in this column. Select the secondary UPS to associate with the node from the list (which is based on the Network UPS Devices table).

Procedure to Configure UPS Protection

1. Check **Enable UPS Support**.
2. If desired, check the **low battery message** option.
This requires both Primary and Secondary UPS devices to have low batteries before the notice is sent to initiate a graceful shutdown.
3. If necessary, **add** network UPS devices.
See [Add Network UPS Device](#) below.
4. Select or change the following from the drop-down lists in the **UPS device table**:
 - UPS Type
 - Primary UPS
 - Secondary UPS
5. Click **OK** to finish.

Add Network UPS Device

Devices need to be added to the Network UPS Devices table on the **UPS** page for the nodes to be associated with them.

1. Click the **Add** button to the right of the Network UPS Devices table.

2. At the **Add Network UPS Device** page, enter:
 - IP Address of the device
 - APC User Name (usually the UPS administrator name, default is **apc**)
 - APC Authentication Phrase (found under low battery shutdown configuration in the APC UPS interface; it is **NOT** the Administrator password)
3. Click **Add**.

You are returned to the **UPS** page and the device is shown in the Network UPS Devices table. The table title UPS count is increased by one. Repeat the process for additional devices.

Change Network UPS Device

To change the settings of a network UPS device:

1. Select a **device** in the Network UPS Devices field to change.
2. Click **Change**.
3. Edit any of the **three options** for the device.
4. Click **Change** again.

Any changes you make are applied to all nodes that are currently using this device.

Delete Network UPS Device

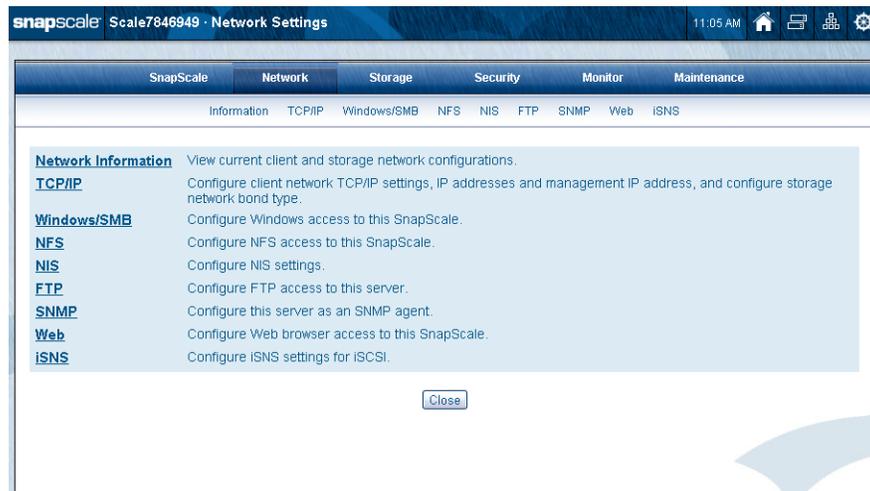
To delete a network UPS device:

1. If the device is still connected to **any nodes**, deselect the device from the nodes.
2. Highlight the **device** in the Network UPS Devices field.
3. Click **Delete**.

The device is deleted from the list.

This section addresses the options for configuring TCP/IP addressing, network bonding, and file access protocols. Network bonding options allow you to configure the SnapScale's Client network for load balancing/failover, Switch Trunking, and Link Aggregation (802.3ad). Network file protocols control how network clients can access the cluster. Access to the cluster's storage space is provided via Windows (SMB), UNIX (NFS), FTP/FTPS, and the Web (HTTP/HTTPS).

NOTE: Uninitialized nodes are configured to use DHCP until they are added to a cluster when they switch to the static IP addresses used by the cluster.¹



Topics in Network Access:

- [View Network Information](#)
- [TCP/IP Networking](#)
- [Windows/SMB Networking](#)
- [NFS Access](#)
- [NIS Domains](#)
- [FTP/FTPS Access](#)
- [SNMP Configuration](#)
- [Web Access](#)
- [iSNS Configuration](#)

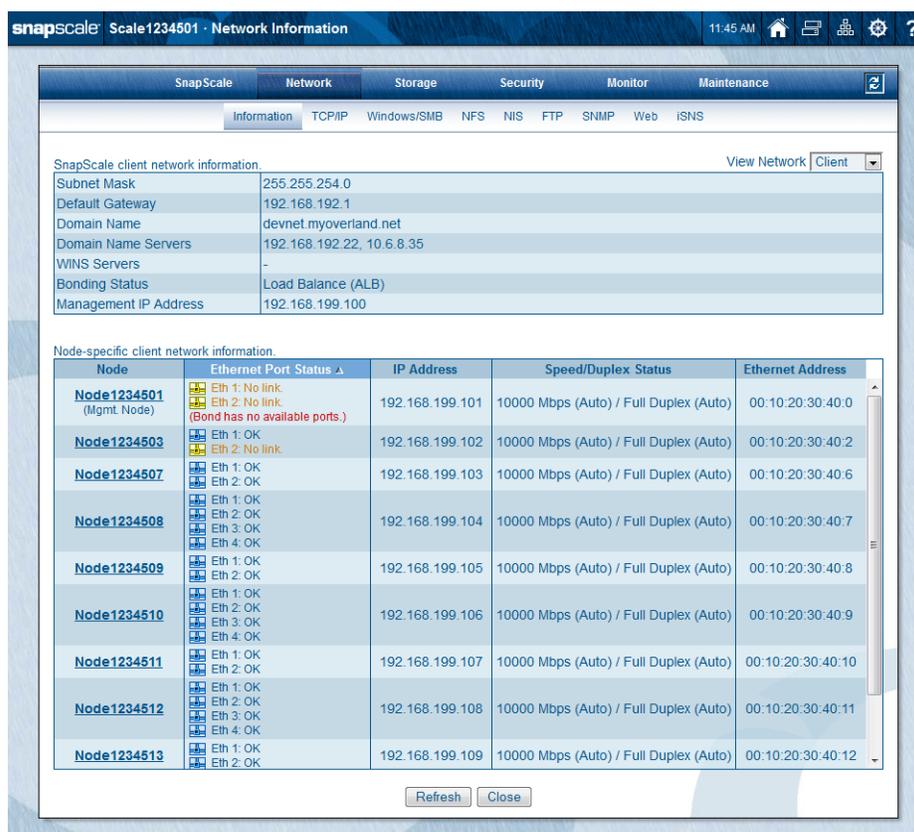
 **IMPORTANT:** The default settings enable access to the SnapScale cluster via all protocols supported by the cluster. As a security measure, disable any protocols not in use. For example, if NFS access to the SnapScale is not needed, disable the protocol in the Web Management Interface under the **Network** tab.

View Network Information

The **Network Information** page displays either the SnapScale's Client or Storage network settings, and identifies the node currently serving as the management node. The information is broken into two parts displaying the common and node-specific network information. Use the **View Network** drop-down menu on the upper right side to select either the Client or Storage network details. Error messages are also shown in this area.

Client Network Information

This page shows the information on the public **Client** network:



The screenshot shows the SnapScale web interface for 'Scale1234501' under the 'Network Information' section. The 'View Network' dropdown is set to 'Client'. The page displays client network information and a table of node-specific information.

Node	Ethernet Port Status	IP Address	Speed/Duplex Status	Ethernet Address
Node1234501 (Mgmt. Node)	Eth 1: No link Eth 2: No link (Bond has no available ports.)	192.168.199.101	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:0
Node1234503	Eth 1: OK Eth 2: No link	192.168.199.102	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:2
Node1234507	Eth 1: OK Eth 2: OK	192.168.199.103	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:6
Node1234508	Eth 1: OK Eth 2: OK Eth 3: OK Eth 4: OK	192.168.199.104	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:7
Node1234509	Eth 1: OK Eth 2: OK	192.168.199.105	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:8
Node1234510	Eth 1: OK Eth 2: OK Eth 3: OK Eth 4: OK	192.168.199.106	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:9
Node1234511	Eth 1: OK Eth 2: OK	192.168.199.107	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:10
Node1234512	Eth 1: OK Eth 2: OK Eth 3: OK Eth 4: OK	192.168.199.108	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:11
Node1234513	Eth 1: OK Eth 2: OK	192.168.199.109	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:40:12

Field definitions are given in the following table:

SnapScale Client Network Information Section	
Subnet Mask	Combines with the IP address to identify the subnet on which the cluster's Client network interfaces are located.

SnapScale Client Network Information Section	
Default Gateway	The network address of the gateway is the hardware or software that bridges the gap between two otherwise unroutable networks. It allows data to be transferred among computers that are on different subnets.
Domain Name	The ASCII name that identifies the DNS domain name that is added to the cluster name to form the fully-qualified host name of the cluster. Additional space-separated domain names are added to the cluster's domain search suffix list.
Domain Name Servers	The IP address of up to three servers that maintain a mapping of all host names and IP addresses for translating domain names into IP addresses.
WINS Servers	The IP address of up to four Windows Internet Naming Service (WINS) servers which locate network resources in a TCP/IP-based Windows network by automatically configuring and maintaining name and IP address mapping tables.
Bonding Status	Shows Load Balance (ALB), Failover, Switch Trunking, or Link Aggregation (802.3ad) as the selected bonding.
Management IP Address	The IP address configured to access and manage the SnapScale cluster through the Web Management Interface.
Node-specific Client Network Information Section	
Node	The name of the specific node. The node designated as the Management node is so noted.
Ethernet Port Status	Shows abbreviated references of the ethernet ports of the node and their statuses. <ul style="list-style-type: none"> • OK – A blue icon () indicates a healthy connection. • No Link – A yellow icon () indicates no link for that port. • Failed – A red icon () indicates that the port has failed.
IP Address	The unique 32-bit value that identifies the node on a network subnet. This is automatically assigned to each node from the pool of IP addresses configured on the cluster.
Speed/Duplex Status	Speed: Ethernet link speed. Duplex Status: Full-duplex; two-way data flow simultaneously.
Ethernet Address	The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet port (xx:xx:xx:xx:xx:xx).

Storage Network Information

This page shows the information on the private **Storage** network:

The screenshot shows the SnapScale Network Information page for a cluster named 'Scale1234501'. The page is divided into two main sections: 'SnapScale storage network information' and 'Node-specific storage network information'.

SnapScale storage network information:

Subnet Mask	255.255.0.0
Bonding Status	Fallover
Multicast IP Address	233.33.0.0

Node-specific storage network information:

Node	Ethernet Port Status	IP Address	Speed/Duplex Status	Ethernet Address
Node1234501 (Mgmt. Node)	Eth 3: OK Eth 4: OK	192.0.2.1	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:0
Node1234503	Eth 3: OK Eth 4: OK	192.0.2.3	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:2
Node1234507	Eth 3: OK Eth 4: OK	192.0.2.7	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:6
Node1234508	Eth 5: OK Eth 6: OK	192.0.2.8	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:7
Node1234509	Eth 3: OK Eth 4: OK	192.0.2.9	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:8
Node1234510	Eth 5: OK Eth 6: OK	192.0.2.10	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:9
Node1234511	Eth 3: OK Eth 4: OK	192.0.2.11	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:10
Node1234512	Eth 5: OK Eth 6: OK	192.0.2.12	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:11
Node1234513	Eth 3: OK Eth 4: OK	192.0.2.13	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:12
Node1234514	Eth 5: OK Eth 6: OK	192.0.2.14	10000 Mbps (Auto) / Full Duplex (Auto)	00:10:20:30:33:13

Buttons: Refresh, Close

Field definitions are given in the following table:

SnapScale Storage Network Information	
Subnet Mask	Combines with the IP address to identify the subnet on which the cluster's Storage network interfaces are located.
Bonding Status	Shows Load Balance (ALB), Failover, Switch Trunking, or Link Aggregation (802.3ad) as the selected bonding.
Multicast IP Address	Multicast address used for inter-node cluster messaging.
Node-specific Client Network Information Section	
Node	The name of the specific node. The node designated as the Management node is so noted.
Ethernet Port Status	Shows abbreviated references of the ethernet ports of the node and their statuses. <ul style="list-style-type: none"> • OK – A blue icon () indicates a healthy connection. • No Link – A yellow icon () indicates no link for that port. • Failed – A red icon () indicates that the port has failed.
IP Address	The unique 32-bit value that identifies the node on a network subnet. This is automatically assigned to each node from the pool of IP addresses configured on the cluster.
Speed/Duplex Status	Speed: Ethernet link speed. Duplex Status: Full-duplex; two-way data flow simultaneously.

Node-specific Client Network Information Section

Ethernet Address	The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet port (xx:xx:xx:xx:xx:xx).
------------------	---

TCP/IP Networking

SnapScale nodes ship with either four 1GbE or 10GbE ports at the rear for network connections. The Storage network ports are always bonded using Failover mode. The Client network ports are bonded by default using Load Balance (ALB), but can be changed after the cluster is created to one of the other bonding modes:

- Failover
- Switch Trunking
- Link Aggregation (802.3ad)

See [Bonding Options](#) on page 4-5 for descriptions.

The **TCP/IP Networking** page provides configuration of the common cluster network settings, the static Management IP address, and the pool of static IP addresses to automatically assign to cluster nodes.

NOTE: If the Client network runs a DHCP server, be sure the static IP addresses assigned to the nodes and Management IP are excluded from DHCP assignment.

The screenshot displays the SnapScale TCP/IP Networking configuration interface. The top navigation bar includes tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Network' tab is active, and the 'TCP/IP' sub-tab is selected. The main content area is divided into two sections:

SnapScale client network settings:

- Subnet Mask: 255.255.0.0
- WINS Servers: Three optional input fields.
- Default Gateway: 10.25.1.1 (optional)
- DNS Domain Name: devnet.myoverland.net (optional)
- Domain Name Servers: 10.6.8.34 (optional) and 10.6.8.35 (optional)
- Bond Type: Load Balance (ALB)

SnapScale management and node client network static IP addresses:

- Static IP Address:** A list of IP addresses with checkboxes:
 - 10.25.11.160 (Management IP address.)
 - 10.25.11.161 (Node IP address.)
 - 10.25.11.162 (Node IP address.)
- Optional: Enter a starting IP address and click "Populate" to populate the list on the left with sequential IP addresses.
- Starting IP Address: [Input field] [Populate Static IP Addresses] (optional)

At the bottom, there are buttons for OK, Utility IP Address, Storage Network Properties, and Cancel.

The following table describes the configuration options found on the **TCP/IP Networking** page:

Column	Description
Subnet Mask	Combines with the IP address to identify the subnet on which the cluster's Client network interfaces are located.
WINS Servers	The IP address of up to four Windows Internet Naming Service (WINS) servers which locate network resources in a TCP/IP-based Windows network by automatically configuring and maintaining name and IP address mapping tables.
Default Gateway	The network address of the gateway is the hardware or software that bridges the gap between two otherwise unroutable networks. It allows data to be transferred among computers that are on different subnets.
DNS Domain Name	The ASCII name that identifies the DNS domain name that is added to the cluster name to form the fully-qualified host name of the cluster, and also serves as the primary DNS search suffix. Additional space-separated domain names can be specified to extend the domain search suffix list.
Domain Name Servers	The IP address of up to three servers that maintain a mapping of all host names and IP addresses for translating domain names into IP addresses.
Bond Type	Use the drop-down list to select one of the four bonding modes for the Client network interface on all nodes.
Static IP Address	This table shows the SnapScale Management IP address and the pool of Client network static IP addresses to be automatically assigned by the cluster to the different nodes. To change or populate the list with a contiguous range of IP addresses, in the area to the right, enter a starting IP address and click Populate Static IP Addresses .

Bonding Options

The bonding options available for SnapScale nodes:

- **Failover** – This default mode uses one Ethernet port as the primary network interface and one port held in reserve as the backup interface. Redundant network interfaces ensure that an active port is available at all times. If the primary port fails due to a hardware or cable problem, the second port assumes its network identity. The ports on a node should be connected to different switches (though this is not required).

Default ports are:

- Basic 1GbE X2 or X4 – **Port 1**.
- Single or Dual 10GbE Card X2 – **Port 3**.
- Single or Dual 10GbE Card X4 – **Port 5**.

NOTE: Failover mode provides switch fault tolerance, as long as ports are connected to different switches.

- **(Automatic) Load Balance (ALB)** – An intelligent software adaptive agent repeatedly analyzes the traffic flow from the node and distributes the packets based on destination addresses, evenly distributing network traffic for optimal network performance. Both ports of the bond need to be connected to the same switch or logical switch.
- **Switch Trunking** – This mode groups multiple physical Ethernet links to create one logical interface. Provides high fault tolerance and fast performance between switches, routers, and servers. Both ports of the bond need to be connected to the same physical or logical switch, and the switch ports must be configured for static link aggregation.
- **Link Aggregation (802.3ad)** – This method of combining or aggregating multiple network connections in parallel is used to increase throughput beyond what a single connection could handle. It also provides a level of redundancy in case one of the links fails. It uses Link Aggregation Control Protocol (LACP), also called dynamic link aggregation, to autonegotiate trunk settings. Both ports of the bond need to be connected to the same switch or logical switch.

Guidelines in TCP/IP Configuration

Consider the following guidelines when connecting a SnapScale cluster to the network.

Configure the DNS for Name Resolution and Round Robin Load Distribution

To evenly distribute client access loads to the cluster nodes, add a DNS A record for the cluster name for each IP in the node IP address pool. The DNS server then rotates through the node IP addresses in a round-robin basis when serving name resolution requests for the cluster name.

Do not add an A record for the cluster name pointing to the Management IP address. If desired, or if using Snap EDR, add an A record for the cluster name followed by “-MGT” for the Management IP address. For example, if the cluster name is Scale1234567, create an A record for hostname “Scale1234567-MGT.”

Make Sure the Switch is Set to Autonegotiate Speed/Duplex Settings

All Ethernet ports on the cluster nodes are set to autonegotiate speed and duplex settings with the Ethernet switch. The switch to which the SnapScale is connected *must* be set to autonegotiate; otherwise, network throughput or connectivity to the node may be seriously impacted.

Cluster Restart Required when Switching the Storage Network to or from Switch Trunking or Link Aggregation (802.3ad)

To prevent the interruption of communication on the Storage network during reconfiguration of the Storage switch, the cluster must be shutdown before changing the Storage network bond setting to or from Switch Trunking or Link Aggregation (802.3ad). After all the Storage switches have been reconfigured, restart the cluster normally by turning the nodes back on.

Configure the Client Switch for Load Balancing

If you select either Switch Trunking or Link Aggregation (802.3ad) network bonding configuration for the Client network bond, be sure the switch is configured correctly for that bonding method **after** configuring the bond on the node. No switch configuration is required for Adaptive Load Balancing (ALB).

Edit Storage Network Properties

 **IMPORTANT:** Changing the bond type for your SnapScale's storage network may require changes to your network switch.

The bond type for the Storage network of a SnapScale cluster can be changed as needed.

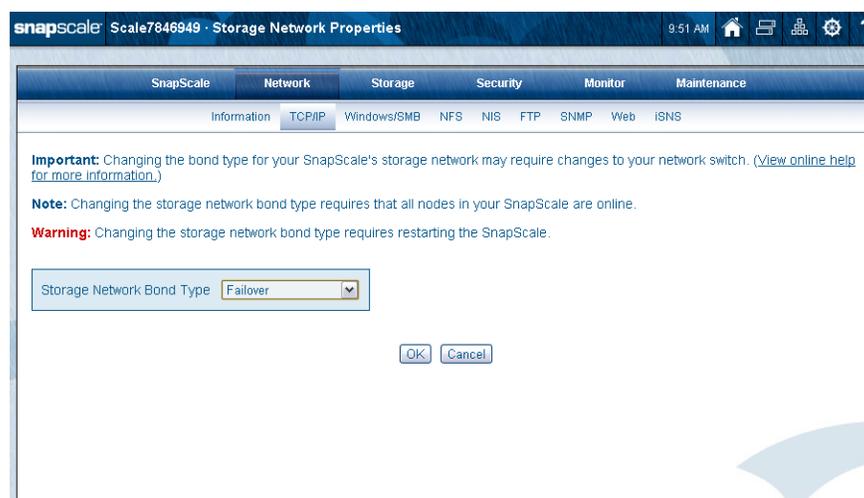
 **CAUTION:** All cluster nodes must be online when their bond type is changed. After changing the bond type, the cluster must be restarted. If the switch is being reconfigured, the cluster must be shut down completely, the Storage network switches reconfigured to the new bond type, and then all nodes restarted.

The following bond types are supported:

- Failover
- Load Balance (ALB)
- Switch Trunking
- Link Aggregation (802.3ad)

See [Client and Storage Networks](#) in [Chapter 1](#) for descriptions.

The following page shows the **bonding options** available from the drop-down list:

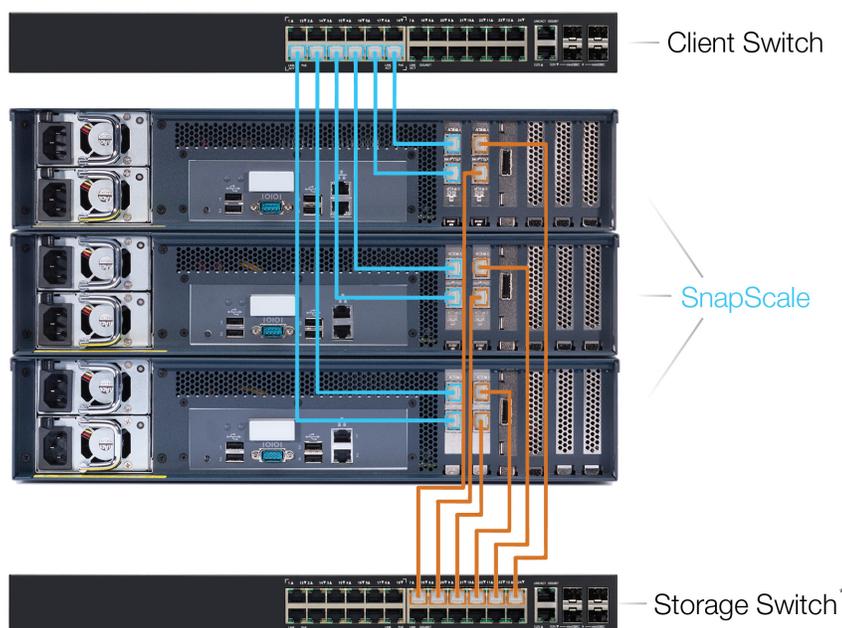


When changing the bond type, depending on the type of change, the following requirements must be met:

- If changing the Storage network between Failover and ALB, the cluster must reboot.
- If changing the Storage network to or from Switch Trunking or Link Aggregation (802.3ad), the cluster must be shut down completely, the Storage network switches reconfigured to the new bond type, and then all nodes restarted.

 **CAUTION:** If you change the bonding mode from the default Failover to ALB, Switch Trunking, or Link Aggregation (802.3ad), you **MUST** re-cable the Storage network ports on each node to the same switch. You **CANNOT** straddle them across two Storage network switches like you do for Failover.

Cabling for ALB, Switch Trunking, or Link Aggregation (802.3ad) Example



Utility IP Address

To assign an additional static IP address to a specific node, click the **Utility IP Address** button on the **TCP/IP Networking** page.

The screenshot shows the SnapScale Utility IP Address configuration page. The page title is "snapscale Scale7846949 · Utility IP Address". The navigation menu includes "SnapScale", "Network", "Storage", "Security", "Monitor", and "Maintenance". Under the "Network" tab, there are sub-tabs for "Information", "TCP/IP", "Windows/SMB", "NFS", "NIS", "FTP", "SNMP", "Web", and "iSNS". The main content area contains the following text:

The utility IP address is a client network static IP address that you specify and assign to a node in the SnapScale. The utility IP address will remain with the node, even if the node is restarted or goes offline.

Important: The utility IP address should **not** be used for client (read/write) access, and must be located on the same subnet as the SnapScale client network (Range: 10.25.0.1 - 10.25.255.254).

Specify a utility IP address and node.

Utility IP Address

Node

OK Cancel

The Utility IP address can be used to reliably access a specific node by a known IP address, and is particularly useful for backup agents and media server installations (see [Appendix A, Backup Solutions](#)). The Utility IP address is assigned to the node in addition to its static IP address automatically assigned from the cluster IP address pool, as well as the Management IP if the node serves as the Management node.

IMPORTANT: The Utility IP address must be located on the same subnet as the SnapScale Client network. The address should be assigned **BEFORE** installing a backup agent or media server on a node. Once the Utility IP address has been assigned, you must add a host record to the DNS server for the node name pointing to the Utility IP address (do **NOT** add it as another host record for the cluster name).

Only one Utility IP address can exist on a cluster. The Web Management Interface will not allow a new Utility IP address to be created if a Utility IP address currently exists, or when an address does not exist but there are one or more offline nodes (which may have an address already configured on them). The Utility IP address also must not be the same as the Management IP address or any existing address in the cluster IP address pool.

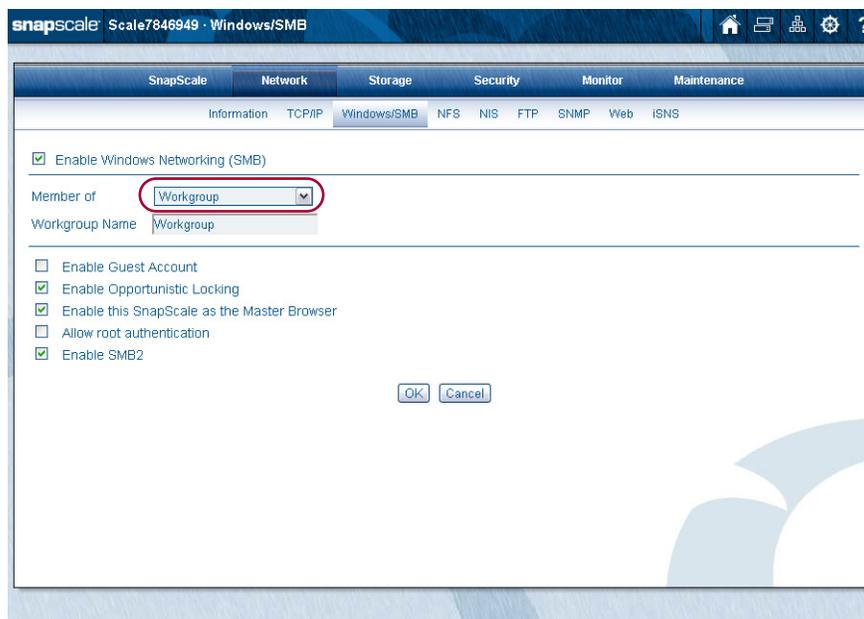
To Configure a Utility IP Address:

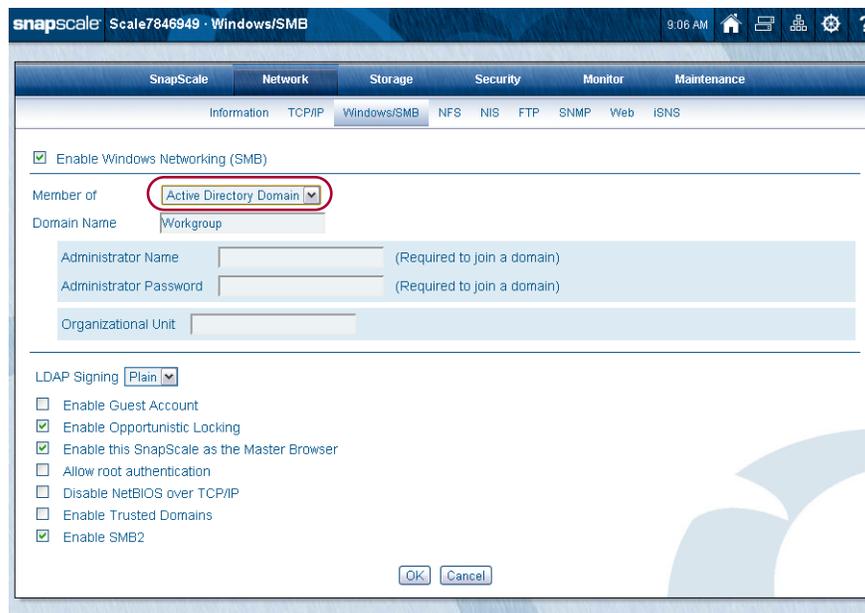
1. On the **Utility IP Address** page, in the empty field, enter a **static IP address** on the same subnet as the IP address pool on the Client network.
2. Using the drop-down list, select the **cluster node** to which the Utility IP address will be assigned.
3. Click **OK**.
4. At the confirmation page, click **Save Changes**.

The Utility address is displayed on the **Network Information** page (**Network > Information**) beneath the static address of the node on which it was configured. The Utility IP address remains with the node, even if the node is restarted or goes offline.

Windows/SMB Networking

Windows/SMB and security settings are configured on the **Windows/SMB** page of the Web Management Interface. You can configure the cluster as a member of either a **Workgroup** or an **Active Directory Domain**, as shown below:





If you run Windows networking in domain mode, you must not configure Date/Time to synchronize with an NTP server.

Support for Windows/SMB Networking

The default settings make the SnapScale available to SMB clients in the workgroup named *Workgroup*. Opportunistic locking is enabled, as is participation in master browser elections.

Consider the following when configuring access for your Windows networking clients.

Support for Microsoft Name Resolution Servers

The SnapScale supports NetBIOS, WINS, and DNS name resolution services. However, when you use Windows Active Directory Services (ADS), make sure forward and reverse name lookups are correctly set up.

ShareName\$ Support

RAINcloudOS supports appending the dollar-sign character (\$) to the name of a share in order to hide the share from SMB clients accessing the SnapScale.

NOTE: As with Windows servers, shares ending in '\$' are not truly hidden, but rather are filtered out by the Windows client. As a result, some clients and protocols can still see these shares.

To completely hide shares from visibility from any protocols, the **Shares** page (**Security > Shares**) provides access to a special share option that hides a share from SMB and HTTP/HTTPS clients. However, shares are not hidden from NFS clients, which cannot connect to shares that are not visible. To hide shares from NFS clients, consider disabling NFS access on hidden shares.

For new shares, select **Create Share** and click the **Advanced Share Properties** button to access the Hidden share option. For existing shares, select the share, click **Properties**, and click **Advanced Share Properties** to access the Hidden share option.

Support for Windows Network Authentication

This section summarizes important facts regarding the RAINcloudOS implementation of Windows network authentication.

NOTE: When a SnapScale cluster joins a domain, it does so under its cluster name (Scalennnnnn). When a domain user is authenticated on a node, the cluster name is used. As such, a user can use any node of the cluster to be authenticated and log on.

Windows Networking Options

Windows environments operate in either workgroup mode, where each SnapScale cluster contains a list of local users it authenticates on its own, or ADS domain mode, where domain controllers centrally authenticate users for all domain members.

Option	Description
Workgroup	In a workgroup environment, users and groups are stored and managed separately on each server or cluster in the workgroup.
Active Directory Service (ADS)	<p>When operating in a Windows ADS domain environment, the SnapScale is a member of the domain and the domain controller is the repository of all account information. Client machines are also members of the domain and users log into the domain through their Windows-based client machines. ADS domains resolve user authentication and group membership through the domain controller.</p> <p>Once joined to a Windows ADS domain, the SnapScale can authenticate SMB users against the domain and can configure share access for domain users. Thus, you must use the domain controller to make modifications to user or group accounts. Changes you make on the domain controller appear automatically on the SnapScale.</p> <p>NOTE: Windows 2000 domain controllers must run SP2 or later.</p>

Kerberos Authentication

Kerberos is a secure method for authenticating a request for a service in a network. Kerberos lets a user request an encrypted “ticket” from an authentication process that can then be used to request a service from a server or cluster. The user credentials are always encrypted before they are transmitted over the network.

The SnapScale supports the Microsoft Windows implementation of Kerberos. In Windows ADS, the domain controller is also the directory server, the Kerberos Key Distribution Center (KDC), and the origin of group policies that are applied to the domain.

NOTE: Kerberos requires the cluster's time to be closely synchronized to the domain controller's time. This means that (1) the cluster automatically synchronizes its time to the domain controller's and (2) NTP cannot be enabled when joined to an ADS domain.

Interoperability with Active Directory Authentication

The SnapScale supports the Microsoft Windows 2000/2003/2008 family of servers that run in ADS mode. any SnapScale can join Active Directory Service domains as a member server. References to the SnapScale's shares can be added to organizational units (OU) as shared folder objects.

NOTE: Windows 2000 domain controllers must run SP2 or later.

Guest Account Access to the SnapScale

The **Windows/SMB** page in the Web Management Interface contains an option that allows unknown users to access the SnapScale using the guest account.

Connect from a Windows Client

Windows clients can connect to the SnapScale using either the cluster name or any IP address in the node IP address pool. However, if possible, clients should use the cluster name to benefit from round robin DNS resolution (see [Configure the DNS for Name Resolution and Round Robin Load Distribution](#) on page 4-7).

To navigate to the cluster using Windows Explorer, use one of these procedures:

- For Microsoft Windows Vista, 2008, and 7 clients, navigate to **Network** > *server_name*.
- For Microsoft Windows XP, 2000, or 2003 clients, navigate to **My Network Places** > *workgroup_name* > *server_name*.

Connect a Mac OS X Client Using SMB

Mac OS X clients can connect using SMB. Specify the cluster name (or an IP address from the node IP address pool) in the Connect to Server window (from **Finder** press **Cmd + K**, or select **Finder** > **Go** > **Connect to Server**) as one of the following:

NOTE: If possible, clients should use the cluster name to benefit from round robin DNS resolution (see [Configure the DNS for Name Resolution and Round Robin Load Distribution](#) on page 4-7).

- `smb://cluster_name`
- `smb://node_ip_address`

Tip: To disconnect from the SnapScale cluster, drag its icon into the Trash.

You can also browse the clusters in the Finder file window, under the Shared tab.

Configure Windows/SMB Networking

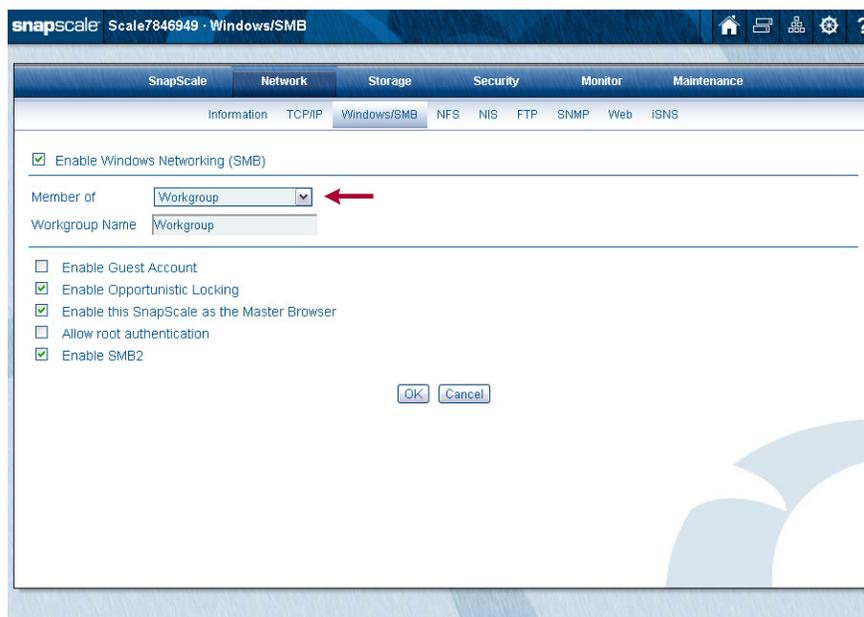
Windows SMB and security settings are configured from this page. The cluster can be configured as part of a Workgroup or an Active Directory Domain.

Before performing the configuration procedures provided here, be sure you are familiar with the information provided previously in [Support for Windows/SMB Networking](#) and [Support for Windows Network Authentication](#) on page 4-12.

To Join a Workgroup

1. Go to **Network** > **Windows/SMB**.

- At the Member list, verify that the default **Workgroup** is selected.



- Edit the **fields** shown in the following table:

Option	Settings
Enable Windows SMB	Check the box to enable SMB and activate the options. Clear the box to disable.
Member Of	Verify that it is set to Workgroup . NOTE: For the Active Directory Domain option, see the following To Join an Active Directory Domain on page 4-15 .
Workgroup Name	The default settings make the SnapScale available in the workgroup named <i>Workgroup</i> . Enter the workgroup name to which the cluster belongs.
Enable Guest Account	Check the box to allow unknown users or users explicitly logging in as Guest to access the SnapScale using the guest account. Clear the box to disable this feature.
Enable Opportunistic Locking	Enabled by default. Opportunistic locking can help performance if the current user has exclusive access to a file. Clear the box to disable this feature.
Allow Root Authentication	Check the box to allow root login to the cluster; clear the box to disable this feature. NOTE: The root password is synchronized with the cluster's admin password.
Enable SMB2	Enabled by default. This more robust version of SMB reduces protocol overhead and is used by default by Windows Vista and later clients. Clear the box to disable this feature (clients that default to SMB2 will automatically connect via SMB1).

- Click **OK** to update Windows network settings immediately.

To Join an Active Directory Domain

When the cluster joins a domain, it does so as a single unit under the cluster name, and all nodes operate equally under the cluster name to authenticate against the domain. This provides multipoint access to the domain through each node.

1. Go to **Network > Windows/SMB**.
2. From the drop-down Member list, select **Active Directory Domain** to view the configuration page.

The screenshot shows the SnapScale Windows/SMB configuration interface. At the top, there are tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Under the Network tab, there are sub-tabs for Information, TCP/IP, Windows/SMB, NFS, NIS, FTP, SNMP, Web, and ISNS. The Windows/SMB configuration page is displayed, showing the following options:

- Enable Windows Networking (SMB)
- Member of: **Active Directory Domain** (indicated by a red arrow)
- Domain Name: Workgroup
- Administrator Name: (Required to join a domain)
- Administrator Password: (Required to join a domain)
- Organizational Unit: (blank)
- LDAP Signing: Plain
- Enable Guest Account
- Enable Opportunistic Locking
- Enable this SnapScale as the Master Browser
- Allow root authentication
- Disable NetBIOS over TCP/IP
- Enable Trusted Domains
- Enable SMB2

Buttons for OK and Cancel are located at the bottom right of the configuration area.

NOTE: You cannot select Active Directory Domain if NTP is enabled.

3. Edit the **fields** shown in the following table:

Option	Description
Enable Windows SMB	Check the box to enable SMB and activate the options. Clear the box to disable.
Member Of	Verify it shows <i>Active Directory Domain</i> .
Domain Name	The default settings make the SnapScale available in the workgroup named <i>Workgroup</i> . Enter the domain name to which the cluster belongs. NOTE: Windows 2000 domain controllers must run SP2 or later.
Administrator Name / Administrator Password	If joining a domain, enter the user name and password of a user with domain join privileges (typically an administrative user).
Organizational Unit	To create a machine account at a different location than the default, enter a name in the field. By default, this field is blank, signaling the domain controller to use a default defined within the controller. NOTE: Sub-organizational units can be specified using Full Distinguished Name LDAP syntax or a simple path ([organizational_unit]/[sub-unit1]/[sub-unit1a])
LDAP Signing	Set LDAP signing for the ADS domain to <i>Plain</i> (no signing), <i>Sign</i> , or <i>Seal</i> , as appropriate for your domain. Default setting is <i>Plain</i> .

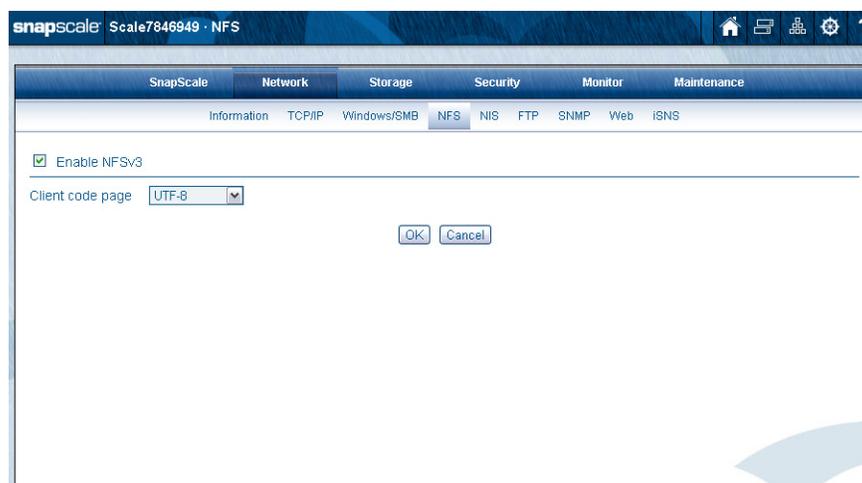
Option	Description
Enable Guest Account	Check the box to allow unknown users or users explicitly logging in as Guest to access the SnapScale using the guest account. Clear the box to disable.
Enable Opportunistic Locking	Enabled by default. Opportunistic locking can help performance if the current user has exclusive access to a file. Clear the box to disable opportunistic locking.
Enable this SnapScale as the Master Browser	Enabled by default. The SnapScale can maintain the master list of all computers belonging to a specific workgroup. (At least one Master Browser must be active per workgroup.) Check the box if you plan to install this cluster in a Windows environment and you want this cluster to be able to serve as the Master Browser for a workgroup. Clear the box to disable this feature.
Allow Root Authentication	Check the box to allow root login to the cluster. NOTE: The root password is synchronized with the cluster's admin password.
Disable NetBIOS over TCP/IP (Active Dir Domain only)	Some administrators may wish to disable NetBIOS over TCP/IP. Check the box to disable NetBIOS; clear the box to leave NetBIOS enabled. NOTE: If you disable NetBIOS and you are joining a domain, you must enter the domain name as a fully qualified domain name (such as, actdirdomainname.companyname.com). A short form such as ActDirDomName does not work.
Enable Trusted Domains (Active Dir Domain only)	SnapScale clusters recognize trust relationships established between the domain to which the SnapScale is joined and other domains in a Windows environment by default. Check the box to enable this feature; clear the box to disable this feature. NOTE: SnapScale clusters remember trusted domains. That is, if this feature is disabled and then activated at a later time, the previously downloaded user and group lists, as well as any security permissions assigned to them, is retained.
Enable SMB2	Enabled by default. This more robust version of SMB reduces protocol overhead and is used by default by Windows Vista and later clients. Clear the box to disable this feature (clients that default to SMB2 will automatically connect via SMB1).

4. Click **OK** to update Windows network settings immediately.

NFS Access

NFS access to the cluster is configured on the **NFS** page of the Web Management Interface. By default, NFS access is enabled and any NFS client can access the SnapScale via NFSv3 with non-root access.

NOTE: NFSv3 is enabled by default. NFSv2 and NFSv4 are not supported.



NFS client access to shares can be specified by navigating to **Security > Shares** and clicking the **NFS Access** link next to the share. You must configure the SnapScale cluster for the code page being used by NFS clients.

Support for NFS

The NFS protocol does not support user-level access control, but rather supports host- and subnet-based access control. On a standard UNIX server, this is configured in an *exports* file. On SnapScale, the exports for each share are configured on the **NFS Access** page independently of user-based share access for other protocols.

SnapScale supports these versions of the NFS protocol and related services:

Protocol	Version	Source
NFS	3.0	RFC 1094, RFC 1813, RFC 3530
Mount	1.0, 2.0, 3.0	RFC 1094 Appendix A, RFC 1813, RFC 3530
Lockd	1.0, 4.0	RFC 1094, RFC1813, RFC 3530

NFS Share Mounting

A share on a SnapScale is equivalent to an exported filesystem on an NFS server. NFS users can mount SnapScale shares, or mount a subdirectory of a share, and access content directly using the following procedure:

1. To mount an **NFS client**, enter the following command:

```
mount cluster_name:/share_name /local_mount
```

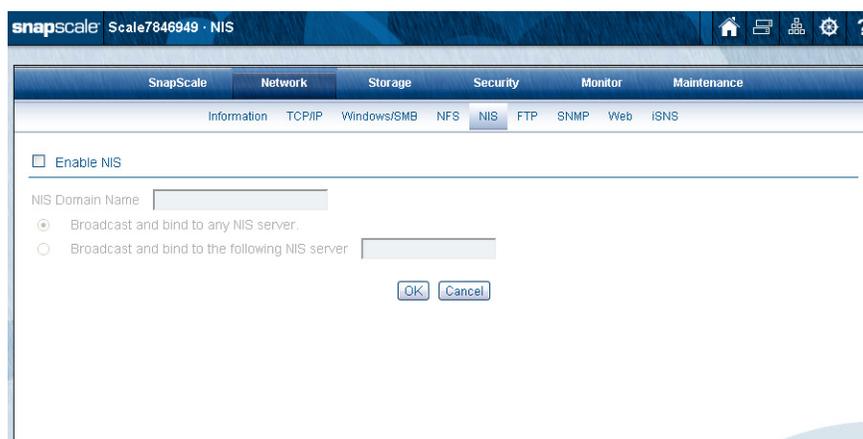
where **cluster_name** is the cluster name (or any address in the node IP address pool), **share_name** is the name of the share you want to mount, and **local_mount** is the name of the mount target directory.

NOTE: If possible, clients should use the cluster name to benefit from round robin DNS resolution (see [Configure the DNS for Name Resolution and Round Robin Load Distribution](#) on page 4-7). Syntax can vary depending upon the operating system.

2. Press **Enter** to connect to the specified share on the cluster.

NIS Domains

NIS domains are configured on the **NIS** page of the Web Management Interface.



The SnapScale cluster can join an NIS domain and function as an NIS client. It can then read the users and groups maintained by the NIS domain. As such, you must use the NIS server to make modifications. Changes you make on the NIS server do not immediately appear on the SnapScale nodes; it may take up to 10 minutes for changes to be replicated.

Guidelines for Configuring NIS

Unless UID/GID assignments are properly handled, NIS users and groups may fail to display properly. For guidelines on integrating compatible SnapScale node UIDs, see [User and Group ID Assignments in Chapter 6](#).

NIS identifies users by UID, not user name, and although it is possible to have duplicate user names, Overland Storage does not support this configuration.

To Join an NIS Domain

1. Go to **Network > NIS**.
2. Edit the **settings** shown in the following table:

Options	Description
Enable NIS	Check the box to enable NIS.
NIS Domain Name	Enter the NIS domain name.
NIS Server	To bind to an NIS server, select either: <ul style="list-style-type: none"> • Broadcast and Bind to Any NIS server to bind to any available NIS servers. • Broadcast and Bind to the following NIS server to bind to a specific NIS server. Enter the NIS server IP address in the field provided.

3. Click **OK** to update the settings immediately.

FTP/FTPS Access

FTP and FTPS settings are configured on the **FTP** page (**Network > FTP**) of the Web Management Interface. FTPS adds encryption to FTP for increased security.



By default, FTP and FTPS clients can access the cluster using the anonymous user account, which is mapped to the SnapScale cluster's *guest* user account and *AllUsers* group account. You can set share access and file access for anonymous FTP users by modifying permissions for these accounts. For more granular control over FTP access, you must create local user accounts for FTP users.

SnapScale also supports explicit FTPS (such as, FTPES or Auth TLS).

NOTE: If standard FTP is enabled, only the data channel is encrypted for FTPS connections – the control channel (including user password) is not encrypted. To force FTPS to encrypt the control channel as well, disable standard FTP.

Supported FTP Clients

SnapScale clusters have been tested with the most common FTP clients and work as expected based on the commands required by RFC 959. SnapScale clusters have been proven to work with these products for standard FTP.

NOTE: Most standard FTP clients do not support FTPS. A client designed to support FTPS is required for FTPS connections.

To Configure FTP/FTPS Access

1. Go to **Network > FTP**.
2. Edit the **settings** shown in the following table:

Option	Settings
Enable FTP	Check the box to enable standard FTP services; leave the box blank to disable access to this cluster via standard FTP.
Enable FTPS	Check the box to enable FTPS services; leave the box blank to disable access to this cluster via FTPS.

Option	Settings
Allow Anonymous User Access	<p>When you allow anonymous login, FTP/FTPS users employ an email address as the password. When you disallow anonymous login, only FTP/FTPS users who are configured as local SnapScale users can access the cluster.</p> <ul style="list-style-type: none"> • Check the box to allow users to connect to the cluster using the anonymous user account. The anonymous user is mapped to the cluster's local guest user account. You can set share access for anonymous FTP/FTPS users by granting either read-write (the default access) or read-only access to the guest account on a share-by-share basis. • Leave the box blank so users cannot log in anonymously but must instead log in via a locally created user name and password.

3. Click **OK** to update the settings immediately.

To Connect via FTP/FTPS

1. To connect to the cluster:
 - For **standard FTP**, enter the cluster's name or IP address in the FTP Location or Address box of a web browser or FTP client application.
 - To connect via a command line, enter:
`ftp cluster_name`
 - To connect via a Web browser, enter:
`ftp://cluster_name`
(where *cluster_name* is the name or IP address of the cluster)
 - For **secure FTPS**, configure your FTPS client application to use explicit FTPS (such as, FTPES or "Auth TLS") and enter the cluster's name or IP address.

NOTE: With anonymous login enabled, access to folders is determined by the share access settings for the guest account. With anonymous login disabled, log into the cluster using a valid local user name and password.

2. Press **Enter** to connect to the FTP root directory.
All shares and subdirectories appear as folders.

NOTE: FTP users cannot manage files or folders in the FTP root directory.

SNMP Configuration

The SnapScale can act as an SNMP agent. SNMP managers collect data from agents and generate statistics and other monitoring information for administrators. Agents respond to managers and may also send traps, which are alerts that indicate error conditions. The cluster communicates with SNMP managers in the same community. A community name is a password that authorizes managers and agents to interact. The cluster only responds to managers that configure the same community strings.

SNMP configuration is accessed by navigating to **Network > SNMP**.

The screenshot shows the SnapScale web interface for SNMP configuration. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Network' tab is active, and the 'SNMP' sub-tab is selected. The 'Enable SNMP' checkbox is checked. Below it, there are four input fields: Read-Only Community (snap_public), Read-Write Community (snap_private), Location (location), and Contact (root@localhost). The 'Enable SNMP traps' checkbox is unchecked, and its associated fields are hidden.

SNMP trap options are hidden until the **Enable SNMP Traps** option is selected.

The screenshot shows the SnapScale web interface for SNMP configuration. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Network' tab is active, and the 'SNMP' sub-tab is selected. The 'Enable SNMP' checkbox is checked. Below it, there are four input fields: Read-Only Community (snap_public), Read-Write Community (snap_private), Location (location), and Contact (root@localhost). The 'Enable SNMP traps' checkbox is checked, and its associated fields are now visible: IP Address 1 (0.0.0.0), IP Address 2 (0.0.0.0), IP Address 3 (0.0.0.0), and IP Address 4 (0.0.0.0). The 'Send a test trap to listed IP addresses upon saving settings.' checkbox is unchecked.

Default Traps

A *trap* is a signal from the SnapScale cluster or any individual node informing an SNMP manager program that an event has occurred. RAINcloudOS supports the default traps shown in this table:

Trap	Initiating Action
coldStart	Whenever SNMP is enabled and a node boots.

Trap	Initiating Action
linkDown	A node's Ethernet interface has gone offline.
linkUp	A node's Ethernet interface has come back online.
authenticationFailure	An attempt to query the SNMP agent using an incorrect read-only or read-write community string was made, and resulted in a failure.
enterpriseSpecific	<p>SnapScale-generated traps that correspond to the error-level, warning-level, and fatal-error-level traps of RAINcloudOS. These traps contain a descriptive message that helps to diagnose a problem using the following OID's:</p> <ul style="list-style-type: none"> 1.3.6.1.4.1.6411.2000.1000.1:loglevel 0 syslog messages (<i>emergency</i>) 1.3.6.1.4.1.6411.2000.1001.1:loglevel 1 syslog messages (<i>alert</i>) 1.3.6.1.4.1.6411.2000.1002.1:loglevel 2 syslog messages (<i>critical</i>) 1.3.6.1.4.1.6411.2000.1003.1:loglevel 3 syslog messages (<i>error</i>) <p>NOTE: There is no specific MIB that defines traps sent by SnapScale clusters or nodes.</p>

Supported Network Manager Applications and MIBs

SnapScale clusters respond to requests for information in MIB-II (RFC 1213) and the Host Resources MIB (RFC 2790 or 1514). You can use any network manager application that adheres to the SNMP V2 protocol with the SnapScale. The following products have been successfully tested with SnapScale clusters: CA Unicenter TNg, HP Open View, and Tivoli NetView.

Configure SNMP

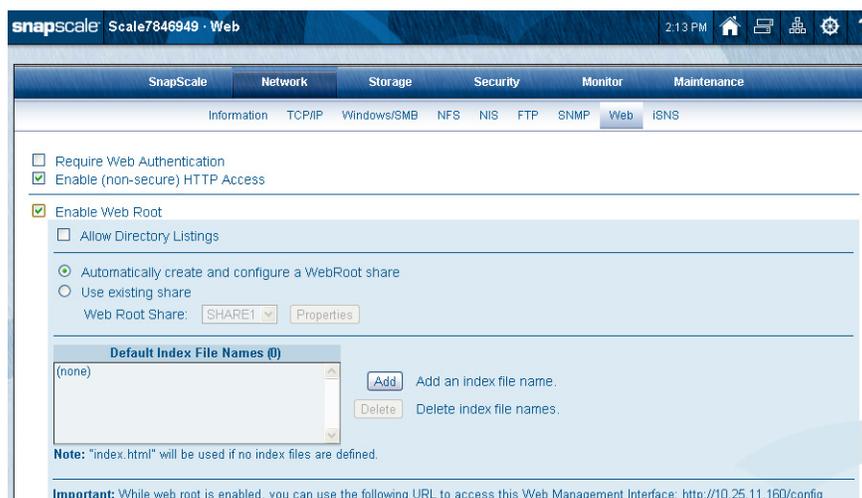
1. Navigate to **Network > SNMP**.
2. Check the **Enable SNMP** box.
3. Edit the settings as described in the following table, and then click **OK**. Once enabled, SNMP managers can access MIB-II and Host Resources MIBs management data on the cluster.

Option	Description
Read-Only Community	<p>To enable SNMP managers to read data from this cluster, enter a read-only community string or accept the default <i>snap_public</i>.</p> <p>NOTE: As a precaution against unauthorized access, Overland Storage recommends that you create your own community string.</p>
Read-Write Community	<p>A read-write string is used for compatibility purposes. Enter a read-write community string or accept the default <i>snap_private</i>.</p> <p>NOTE: As a precaution against unauthorized access, Overland Storage recommends that you create your own community string.</p>

Option	Description
Location	Enter information that helps a user identify the physical location of the cluster nodes. For example, you might include a street address for a small business, a room location such as <i>Floor 37, Room 308</i> , or a position in a rack, such as <i>rack slot 12</i> .
Contact	Enter information that helps a user report problems with the cluster. For example, you might include the name and title of the system administrator, a telephone number, pager number, or email address.
Enable SNMP Traps	Check the box to enable traps. Clear the box to disable SNMP traps.
IP Address 1-4 (only when SNMP Traps are enabled)	Enter the IP address of at least one SNMP manager in the first field as a trap destination. You can enter up to three additional IP addresses.
Send a Test Trap (only when SNMP Traps are enabled)	To verify your settings, check the box. A test message is sent when you click OK.

Web Access

HTTP and HTTPS are used for browser-based access to the cluster via Web View, Web Root, or the Web Management Interface. HTTPS enhances security by encrypting communications between client and cluster, and cannot be disabled. You can, however, disable HTTP access on this **Web** page. Additionally, you can require browser-based clients to authenticate to the cluster.



Configuring HTTP/HTTPS

You can require web authentication, disable HTTP (non-secure) access, and enable the Web Root feature. All HTTP access is made via the root node and the Management IP address.

To Require Web Authentication

Edit the following option and click **OK**.

Option	Description
Require Web Authentication	Check the Require Web Authentication box to require clients to enter a valid user name and password in order to access the cluster via HTTP/HTTPS. Leave the box blank to allow all HTTP/HTTPS clients access to the cluster without authentication. NOTE: This option applies to both Web View and Web Root modes.

To Enable HTTP Access to the SnapScale Cluster

Edit the following option and click **OK**.

Option	Description
Enable (non-secure) HTTP Access	Check the Enable HTTP Access box to enable non-secure HTTP access. Leave the box blank to disable access to the cluster via HTTP. NOTE: This option applies to both Web View and Web Root modes.

To Connect via HTTPS or HTTP

1. Enter the **cluster name**, Management IP address, or any IP address from the node IP address pool in a Web browser.

Web access is case-sensitive. Capitalization must match exactly for a Web user to gain access. To access a specific share directly, Internet users can append the full path to the SnapScale name or URL, as shown in the following examples:

```
https://Node2302216/Share1/my_files
```

```
https://10.10.5.23/Share1/my_files
```

2. Press **Enter**.

The **Web View** page opens.

Using Web Root to Configure the SnapScale as a Simple Web Server

When you enable the Web Root feature from the **Web** page, you can configure your SnapScale cluster to open automatically to an HTML page of your choice when a user enters the following in the browser field:

```
http://[cluster_name] or http://[IP address]
```

In addition, files and directories underneath the directory you specify as the Web Root can be accessed by reference relative to `http://[cluster_name]` without having to reference a specific share. For example, if the Web Root points to the directory *WebRoot* on share *SHARE1*, the file *SHARE1/WebRoot/photos/slideshow.html* can be accessed from a web browser:

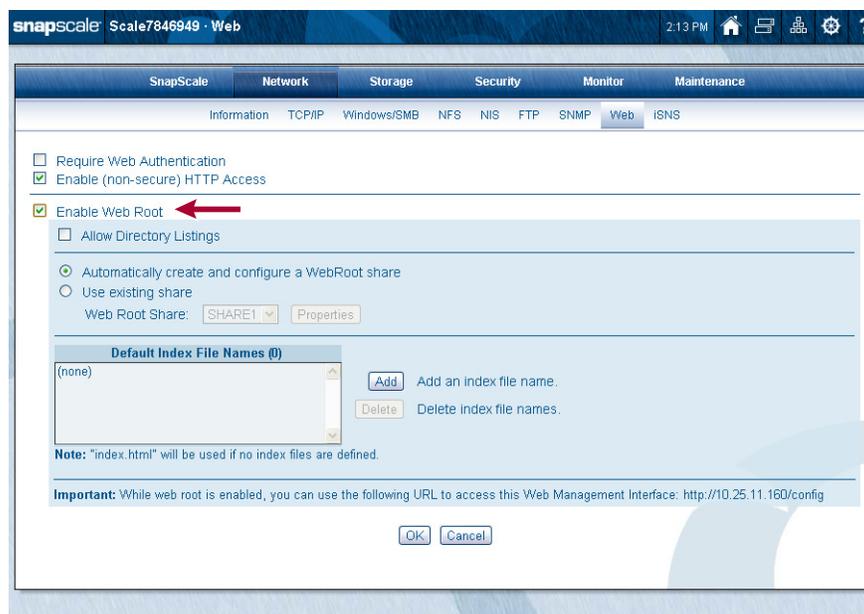
```
http://[cluster_name]/photos/slideshow.html
```

The Web Root can also be configured to support directory browsing independent of Web View (access through shares).

NOTE: SnapScale supports direct read-only web access to files. It is not intended for use as an all-purpose Web Server, as it does not support PERL or Java scripting, animations, streaming video, or anything that would require a special application or service running on the SnapScale cluster.

Configuring Web Root

Check the **Enable Web Root** box to configure the SnapScale to serve the Web Root directory as the top level web access to the SnapScale cluster, and optionally, automatically serve an HTML file inside. When the box is checked, the options described below appear.



1. Complete the following information, then click **OK**.

Option	Description
Allow Directory Listings	<p>If Allow Directory Listings is checked and no user-defined index pages are configured or present, the browser opens to a page allowing browsing of all directories underneath the Web Root.</p> <p>NOTE: Checking or unchecking this option only affects directory browsing in Web Root. It does not affect access to Web View directory browsing.</p>
Create and configure a Web Root share	<p>Select one of the following:</p> <ul style="list-style-type: none"> • Automatically create and configure a Web Root share: A share named "WebRoot" is automatically created. By default, the share is hidden from network browsing and has all network access protocols except HTTP/HTTPS enabled (as such, it can be accessed from a browser as the Web Root but can not be accessed via Web View). You can change these settings at Security > Shares. • Use existing share: From the drop-down list of existing shares for selection, select a share and click the Properties button to edit the selected share's properties (see Security > Shares).

Option	Description
Default Index File Names	<p>Files found underneath the Web Root with names matching those in this list is automatically served to the web browser when present, according to their order in the list. To add a filename, click the Add button, enter the name of one or more index HTML files, then click OK. The file you entered is shown in the Index Files box.</p> <p>NOTE: If no files are specified, <code>index.html</code> is automatically used if found.</p> <p>To delete a name, highlight it and click Delete. At the confirmation page, click Delete again.</p>

2. Map a drive to the **share** you have designated as the Web Root share and upload your HTML files to the root of the directory, making sure the file names of the HTML files are listed in the Index Files box.

Accessing the Web Management Interface when Web Root is Enabled

By default, when you connect to a SnapScale cluster with Web Root enabled, the browser loads the user-defined HTML page or present a directory listing of the Web Root. To access the Web Management Interface (for example, to perform administrative functions or change a password), enter the following in the browser address field:

```
http://[nodename or ip address]/config
```

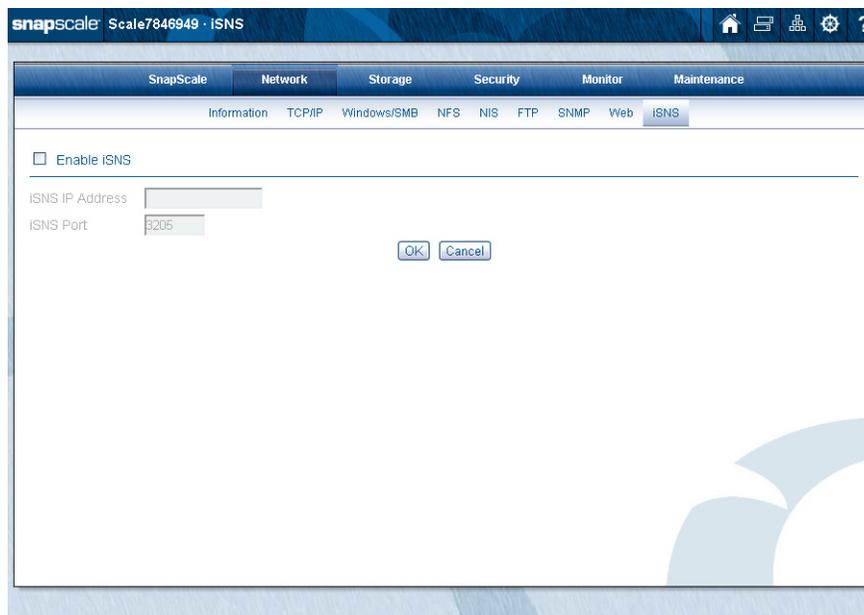
You are prompted for your User ID and password, then you are placed into the Web Management Interface.

If you need to access the **Web View** page to browse shares on the cluster independent of Web Root, enter this in the browser address:

```
http://[nodename or ip address]/sadmin/GetWebHome.event
```

iSNS Configuration

Microsoft iSNS Server can be used for the discovery of SnapScale iSCSI targets on an iSCSI network.

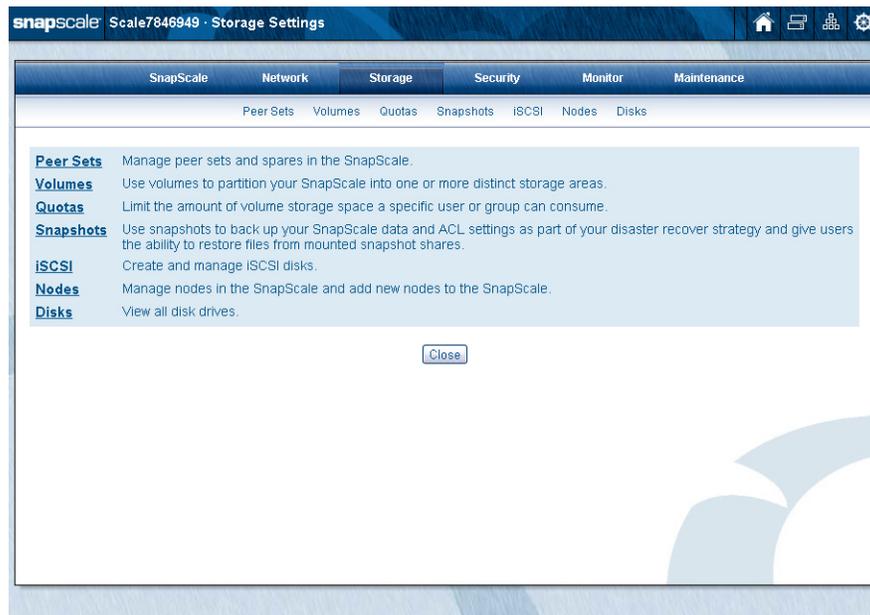


To configure the iSNS settings:

1. If not already installed, install the iSNS service on a **Windows server**.
Note the IP address of the server or workstation on which the iSNS service is installed.
2. Configure iSNS on the **SnapScale**.
On the **Network > iSNS** page, check the **Enable iSNS** box, enter the IP address of the iSNS server, and then click **OK**. If the iSNS server does not use the default port, the iSNS port default value of 3205 can be changed on this page as well.
3. Configure the **iSCSI initiator** to discover iSCSI targets via the iSNS server.

NOTE: After you have completed this procedure, all the iSCSI targets on the SnapScale automatically appear in the Microsoft Initiators target list.

From the storage default page (**Storage Settings**), you can access and configure the storage options for your SnapScale cluster including nodes and drives.



Topics in Storage Options:

- [Peer Sets](#)
- [Volumes](#)
- [Quotas](#)
- [Snapshots](#)
- [iSCSI Disks](#)
- [Nodes](#)
- [Disks](#)

Peer Sets

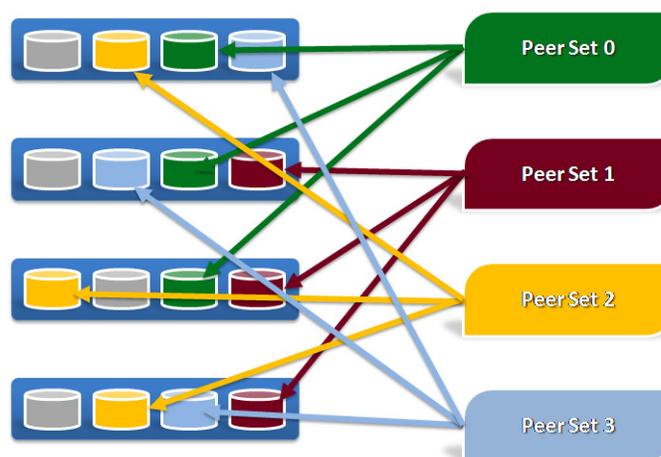
In a cluster, a node is a file server working in tandem with other nodes. The drives on every node are grouped into peer sets or hot spares. Each peer set contains two or three drives, depending on the Data Replication Count, that mirror the same data. To ensure availability, each drive in a peer set resides in a different node.

SnapScale aggregates all the storage on the peer sets in the cluster to form a unified data storage space for network client access. Data access is transparent between the cluster storage space and the peer sets so that users never directly access the peer sets.

When you create a cluster or add new nodes to an existing cluster, SnapScale automatically creates peer sets with the available drives. By distributing peer set members throughout the cluster, the system ensures that content is protected from failure of either individual drives or entire nodes. When they are created, peer sets are assigned a unique peer set ID.

Nodes can be added to expand cluster storage at any time. Based on the configuration settings, the additional drives are either used to create more peer sets or left as hot spares. Nodes can be removed from a cluster for replacement with a new node, and the drives in the replacement node are automatically synchronized with the existing peer sets.

On a four-node cluster configured for 3x replication count, four hot spares, and four drives per node, the peer set formation might look something like this:



Each peer set has members on three different nodes, shown below as peer set 0, 1, 2, and 3. Hot spares are automatically distributed throughout the cluster in order to replace any failed peer set member. When a peer set member fails, a hot spare is assigned from a node on which the peer set does not already have an active member.

The example above uses a 3x Data Replication Count, which means that each peer set contains three members, and as a result all data is replicated three times. The cluster can also be configured for a 2x Data Replication Count, in which case the distribution of two-member peer sets would be different. The system automatically determines which drives are used to form each peer set; you cannot choose them.

NOTE: The Data Replication Count can be decreased from 3x to 2x to increase cluster storage, but cannot be increased from 2x to 3x once the cluster is created.

Peer Sets and Recovery

Though data on peer sets is served indirectly by the unified cluster storage space, access to files stored on a given peer set is dependent on the health of that peer set. When a drive in a peer set fails, data is served from the remaining peer set member drives. If there is a spare reserved for the cluster that does not exist on the same node as another active member of the peer set and is not smaller than other members, the peer set can claim the drive and rebuild the data (using the integrated RapidRebuild feature) onto that spare without administrator intervention.

Peer Set	Status	Member 1	Member 2	DSM
PeerSet16	RapidRebuild: 77% complete.	Node2414532: Disk 12 - 931.51 GB	Node2414528: Disk 2 - 1.82 TB	931.5 GB
PeerSet10	RapidRebuild: 69% complete.	Node2414532: Disk 8 - 931.51 GB	Node2414528: Disk 8 - 931.51 GB	OK
PeerSet4	RapidRebuild: 0% complete.	Node2414532: Disk 4 - 931.51 GB	Node2414528: Disk 4 - 931.51 GB	OK
PeerSet15	RapidRebuild: 0% complete.	Node2414538: Disk 12 - 931.51 GB	Node2414528: Disk 12 - 931.51 GB	OK
PeerSet0	OK	Node2414532: Disk 1 - 931.51 GB	Node2414538: Disk 1 - 931.51 GB	OK
PeerSet1	OK	Node2414532: Disk 2 - 931.51 GB	Node2414538: Disk 2 - 931.51 GB	OK
PeerSet2	OK	Node2414532: Disk 3 - 931.51 GB	Node2414538: Disk 3 - 931.51 GB	OK
PeerSet3	OK	Node2414538: Disk 4 - 931.51 GB	Node2414528: Disk 3 - 931.51 GB	OK
PeerSet5	OK	Node2414532: Disk 5 - 931.51 GB	Node2414538: Disk 5 - 931.51 GB	OK
PeerSet6	OK	Node2414528: Disk 5 - 931.51 GB	Node2414538: Disk 6 - 931.51 GB	OK
PeerSet7	OK	Node2414528: Disk 6 - 931.51 GB	Node2414532: Disk 6 - 931.51 GB	OK
PeerSet8	OK	Node2414532: Disk 7 - 931.51 GB	Node2414538: Disk 7 - 931.51 GB	OK
PeerSet9	OK	Node2414538: Disk 8 - 931.51 GB	Node2414528: Disk 7 - 931.51 GB	OK
PeerSet11	OK	Node2414532: Disk 9 - 931.51 GB	Node2414538: Disk 9 - 931.51 GB	OK
PeerSet12	OK	Node2414528: Disk 9 - 931.51 GB	Node2414538: Disk 10 - 931.51 GB	OK
PeerSet13	OK	Node2414528: Disk 10 - 931.51 GB	Node2414532: Disk 10 - 931.51 GB	OK
PeerSet14	OK	Node2414532: Disk 11 - 931.51 GB	Node2414538: Disk 11 - 931.51 GB	OK

If a peer set is missing one drive but at least one other drive is available, the peer set continues to be accessible but is in degraded mode. This table shows the different peer set statuses:

Peer Set Status	Failure Type	Data Availability
OK	The peer set drives are healthy and connected.	Data is fully available for read and write.
RapidRebuild	Spare made available to rebuild the peer set using RapidRebuild.	Data is fully available for read and write.
Degraded	One drive missing from the peer set	Data is fully available for read and write.
Degraded – Cannot repair; no spares	The peer set cannot be repaired because there are no spare drives.	Data is fully available for read and write.
Degraded – Cannot repair; spares too small	The peer set cannot be repaired because all eligible spares are too small.	Data is fully available for read and write.
Degraded – Cannot repair; spares on same node	The peer set cannot be repaired because the only eligible spares are located on the same node as an active member of the peer set.	Data is fully available for read and write.
Failed	All drives in peer set have failed.	No availability. Contact Overland Support.
Initializing	The peer set is being created or initialized.	Data is not yet available.
Inconsistent	The peer set has more members than the data replication count.	Contact Overland Support.

Peer Set Utilization

Each file's data is spread across multiple peer sets, and the cluster automatically distributes data for different files throughout the peer sets in the cluster. Metadata for files and directories is independently distributed among different peer sets using a hash algorithm for optimum performance and protection.

Peer Set Basics

New drives are initially configured automatically as spare drives. Subsequently, if enough spare drives exist on different nodes to construct new peer sets but still satisfy the spare count setting, the SnapScale automatically creates new peer sets and expand cluster storage space.

Drives in a cluster do not all need to have the same capacity, but drives in a given peer set should have the same capacity or space is wasted on the larger drives.

The following points must be observed in regards to drives used in the cluster:

- The drives in a cluster must all be the same type of drive (such as SAS) and the same rotational speed.
- The storage capacity of a peer set is limited to the smallest capacity drive in the peer set.

In case of peer drive failure, RAINcloudOS continues to serve data reads and writes to that peer set from another member of the peer set as long as the peer set is not offline. If clients are currently using data on the peer set, it continues to operate as-is.

Data Replication Count is an administrator specified, cluster-wide count of the degree of redundancy of data on the cluster. The Data Replication Count can be either 2x or 3x and determines the number of drives (2 or 3) that make up each peer set.

Hot Spares

Each node can have a number of hot spares in the event of a drive failure. The total number of hot spares for the cluster is user selectable. A suggested number of hot spares for various node sizes is provided. If a peer set member drive fails, data from a healthy peer set drive on another node is re-synced onto an available spare on any node that doesn't have another active member of that peer set, and the spare then becomes a member of that peer set.

Drives added to nodes as additions or as replacements to failed drives are automatically configured as spares. If enough spares exist across different nodes to satisfy the Data Replication Count and the spare drives count, the cluster automatically creates a new peer set out of available spare drives.

Snapshot Limitations

- All snapshots are deleted when:
 - New peer sets are automatically created when new drives are installed.
 - One or more new nodes are added to a cluster.
 - If a complete peer set fails.
- A Snapshot may be deleted if:
 - Any peer set member drive runs out of snapshot space.
 - A second member of a peer set (containing *unique* snapshot data) fails, even though the main file system data may still be healthy.

Peer Sets Page



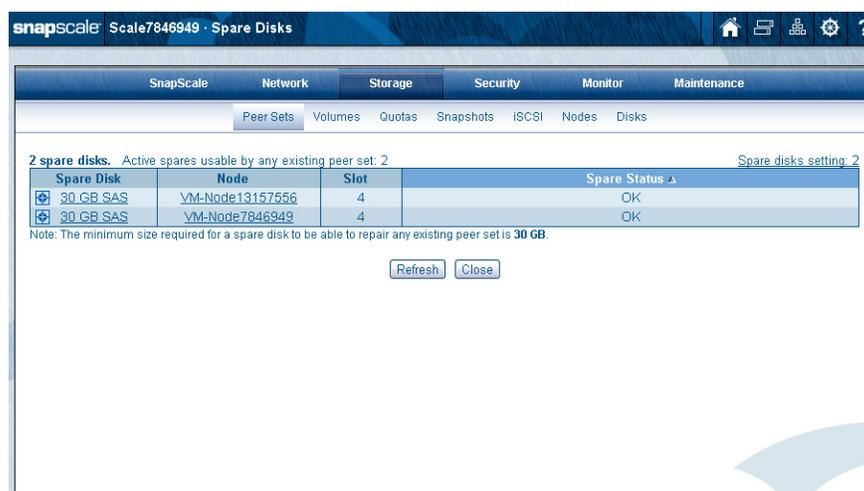
The following table covers the items listed on this page:

Option	Description
# Peer Sets (above table, left)	Displays the total number (#) of peer sets configured and shown in the table.
Data Replication Count (above table, left)	Displays the cluster-wide replication count and links to the SnapScale Properties page.
Active Spare Disks (above table, right)	Shows the number of drives allocated as spares. They are broken out to show the status of spares and the number of spares with that status. <ul style="list-style-type: none"> • OK – A blue icon (OK) indicates spares are active and available. • Too Small – A yellow icon (Too Small) indicates spares are too small to be used in some of the peer sets. A yellow icon with an “X” (Failed) indicates a spare is too small to use with any available peer set. • Failed – A red icon (Failed) indicates spares have failed. Clicking this link opens the Spare Disks page. This is the same as clicking the Spare Disks button at the bottom of the page.
Peer Set	Lists the peer set name and shows a usage bar. Position the cursor over a name (or usage bar) to show the percentage and actual amount of storage space used.
Status	Shows the current status. Refer to Peer Sets and Recovery on page 5-2 for complete details.
Member 1	Shows the node, drive/slot number, and the size of the first member of this peer set. Click to view the Disks page and identify the specific disk drive's location.
Member 2	Shows the node, drive/slot number, and the size of the second member of this peer set. Click to view the Disks page and identify the specific disk drive's location.

Option	Description
Member 3 (if shown)	Shows the node, drive/slot number, and the size of the third member of this peer set when the Data Replication Count is set to "3x." Click to view the Disks page and identify the specific disk drive's location.
DSM (Drive Size Mismatch)	Shows either OK or mismatch size difference. If the member drives are not the exact same size, then capacity is limited to the smallest drive in the peer set, and extra space on larger drives is wasted. In this case, the size displayed reflects the unutilized capacity of the peer set. Position the cursor over a name (or usage bar) to show the unutilized capacity of the peer set.
Spare Disks (button)	Launches the Spare Disks page. See Spare Disks Page below.
Spare Distributor (button)	Launches the Spare Distributor page. See Spare Distributor below.
Data Balancer (button)	Launches the Data Balancer page. See Data Balancer below.
Refresh ( button)	Refreshes the page when clicked.
Close (button)	Closes this page and returns to the Storage Settings page.

Spare Disks Page

When you click the **Spare Disks** button (or the **Active spares** link on the upper right above the table on the **Peer Sets** page), the **Spare Disks** page opens.



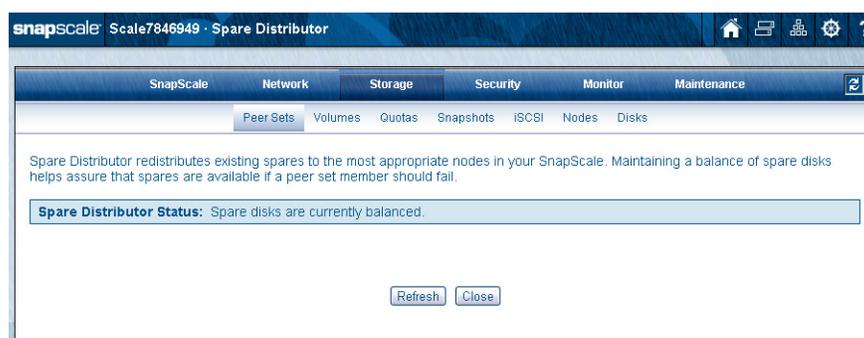
The following table covers the items listed on this page:

Option	Description
# Spare Disks (above table, left)	Displays the total number (#) of spare drives configured and shown in the table.
Active Spares (above table, left)	Displays the cluster-wide number of spares usable by all peer sets.

Option	Description
Spare Disks Setting (<i>above table, right</i>)	Displays the quantity set for spare drives. Clicking this link takes you to the SnapScale Properties page to edit the setting. NOTE: This setting may not equal the number of spare drives currently displayed if there are fewer spare drives available than the setting specifies, or if there is an insufficient number of extra drives to automatically create a new peer set and satisfy the cluster's Data Replication Count.
Spare Disk	Displays disk drive capacity and type. Click a name in the column to open the Disks page and identify a specific disk drive's location.
Node	Displays the name of the node in which the drive resides. Click a name in the column to open the Node Properties page for the specific node.
Slot	Displays the slot number of the listed node where this spare drive is located.
Spare Status	Shows the current status: <ul style="list-style-type: none"> • OK – Spare drive is healthy and can be used by all peer sets. • Spare Too Small – Spare is too small to either repair any existing peer sets or repair <i>n</i> existing peer sets. • Failed – Spare drive has failed.
Refresh ( button)	Refreshes the page when clicked.
Close (<i>button</i>)	Closes this page and returns to the main Peer Sets page.

Spare Distributor

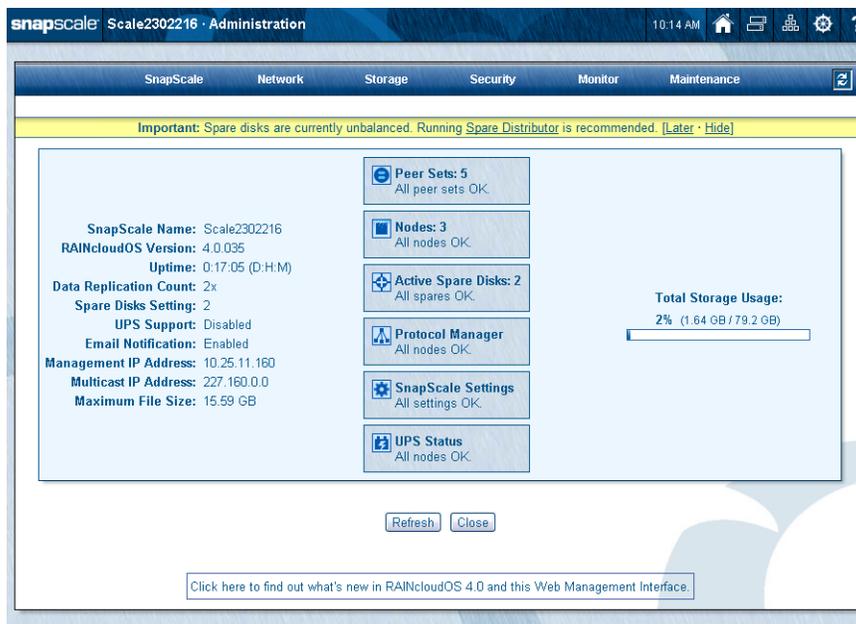
Spare Distributor (formerly the Spare Disk Balancer) evenly redistributes spares and peer set members across the cluster nodes. Maintaining a balance of spare drives helps ensure that spares are available if a peer set member should fail.



Using Spare Distributor

When the cluster detects an uneven distribution of spare drives, an alert banner is displayed in the Web Management Interface and the **Spare Distributor** page is enabled.

NOTE: You can click **Later** to turn off the alert for 24 hours or **Hide** to dismiss the alert.

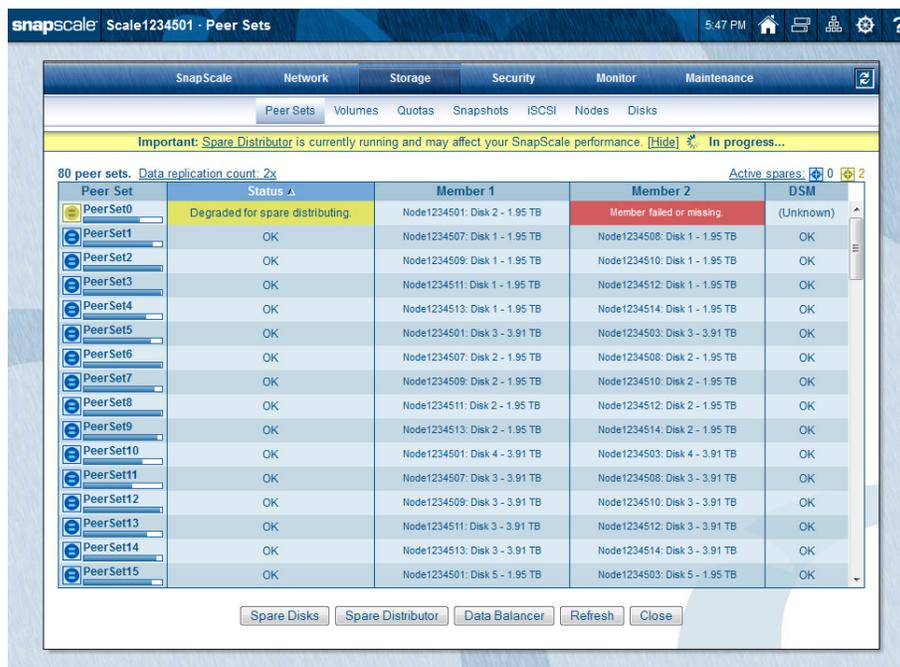


1. Go to **Storage > Peer Sets > Spare Distributor**.

If responding to an alert, you can click the Spare Distributor link in the alert to go directly to the page.

2. Click the **Start Spare Distributor** button to start the process.

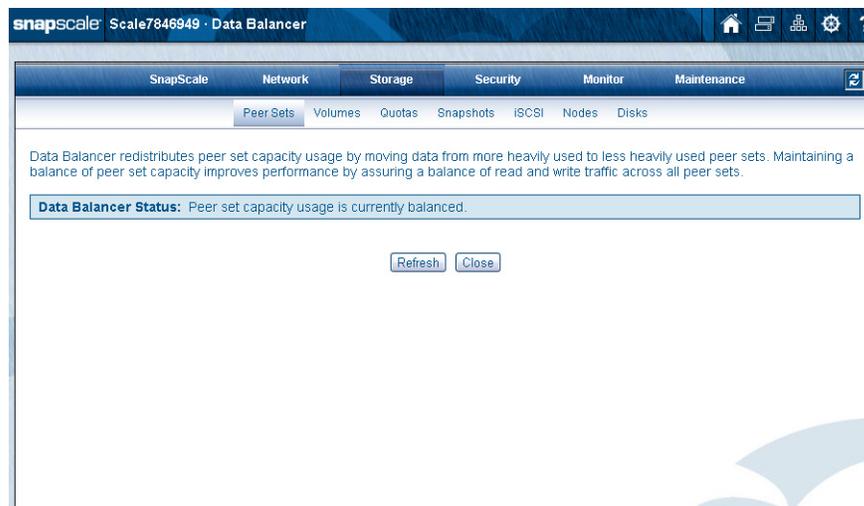
The Spare Distributor redistributes spares and peer set members across the cluster nodes to provide spares on different nodes for better spare availability. Go to the **Peer Sets** page to view the status of the balancer.



If needed, click **Stop Spare Distributor** on the **Spare Distributor** page to stop the operation. Any peer sets currently degraded and being rebuilt by the Spare Distributor will continue with the rebuilding process until completed.

Data Balancer

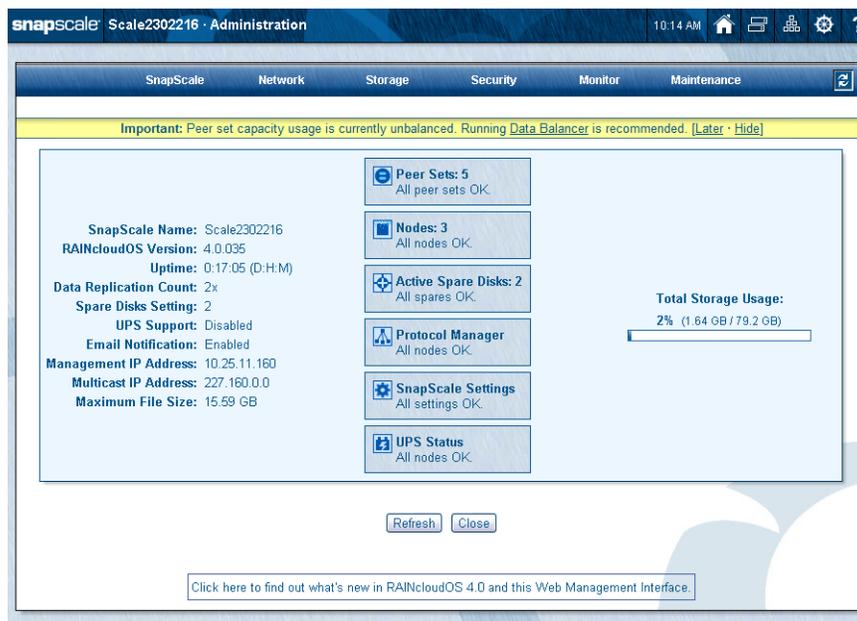
Data Balancer (formerly Capacity Balancer) redistributes peer set utilization by moving data from more to less heavily used peer sets. Maintaining a balance of peer set capacity improves performance by assuring a balance of read and write traffic across all peer sets.



Using Data Balancer

If the peer set data utilization becomes unbalanced, an alert banner is displayed in the Web Management Interface.

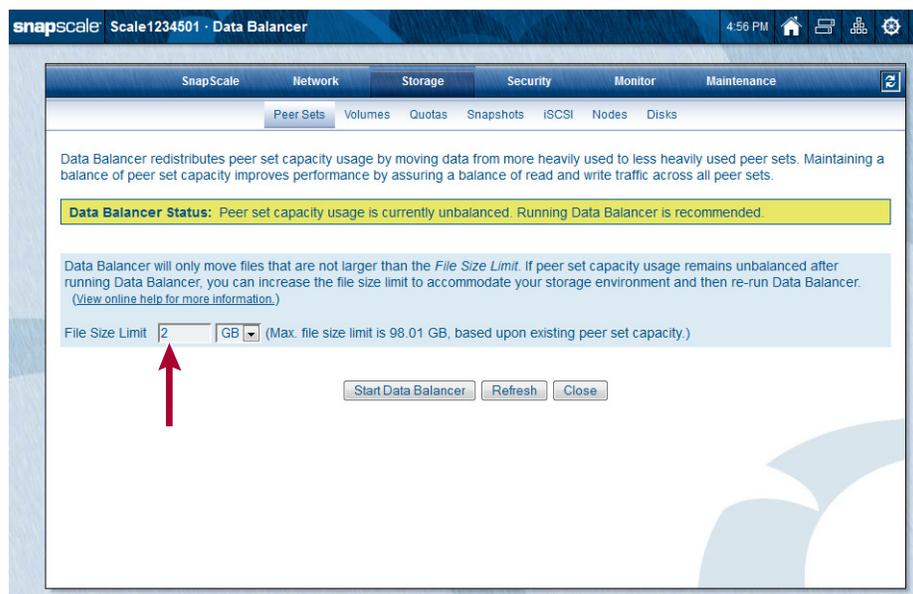
NOTE: You can click Later to turn off the alert for 24 hours or Hide to dismiss the alert.



1. Go to **Storage > Peer Sets > Data Balancer**.

If responding to an alert, you can click the **Data Balancer** link in the alert to go directly to the page.

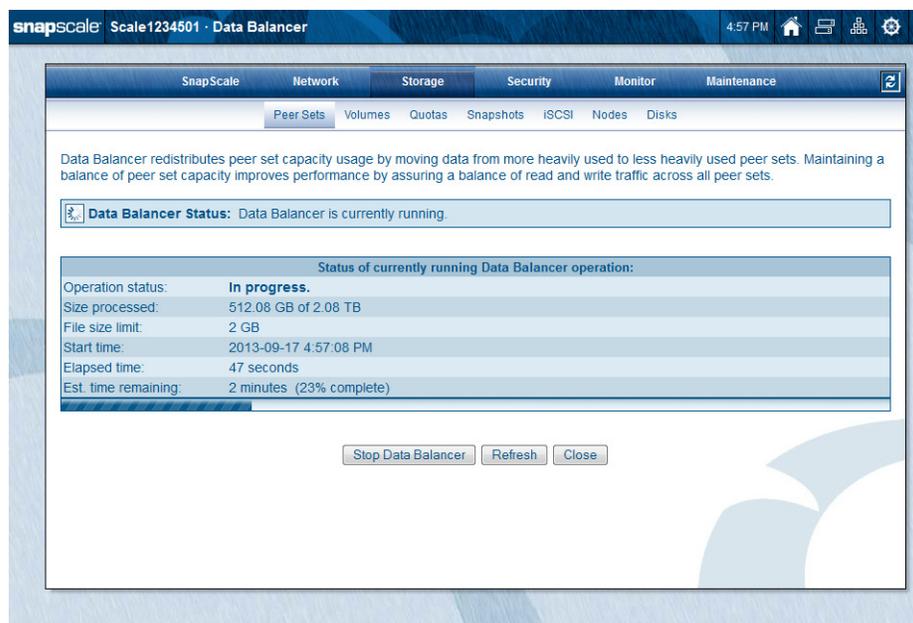
- Review the default **File Size Limit** and change it, if needed.



The **File Size Limit** represents the maximum size of a file the Data Balancer will attempt to move to rebalance peer set consumption. The default is 2GB.

- Click the **Start Data Balancer** button to start the process.

The Data Balancer moves files between peer sets to improve performance and usability. A table is displayed showing that the Data Balancer is running and the percent completed. If needed, click **Stop Data Balancer** to end the operation.

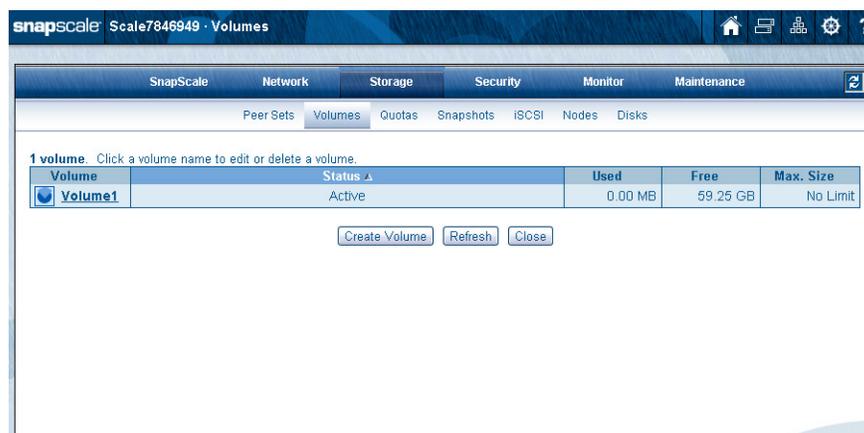


NOTE: The cluster continues to be available for client access during the process. The Data Balancer will skip any file that is currently opened by clients, and will abort moving a file if a client opens it during the move.

An alert banner is displayed on any Administration-level page showing the progress.

Volumes

Use the **Volumes** page (**Storage > Volumes**) to manage the volumes that have been created.



From this page, you can:

- Create a new volume.
- Edit or delete the volume (by clicking the name to access the **Properties** page).

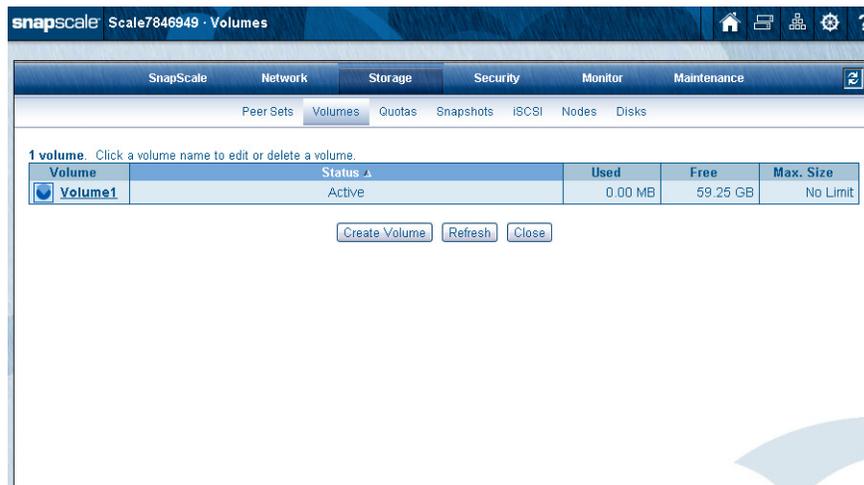
Volume Overview

All the peer sets are unified into a single cluster storage space that can be accessed from any node thus providing multiple access points. One or more volumes can be created to provision the cluster storage:

- All volumes share the same cluster storage space and are thinly provisioned to provide better utilization rates of the space.
- Volumes can be configured with a maximum size setting (quota) to prevent one volume from consuming too much shared cluster storage space. See [Creating Volumes](#) on page 5-12.

Creating Volumes

By default, the full cluster storage space is accessible as one large storage space. However, the storage space can be divided into multiple volumes in order to thinly provision space for specific projects, departments, or roles. Volumes can be constrained to use no more than a certain amount of space available in the clustered storage space. This is done starting at the **Volumes** page:

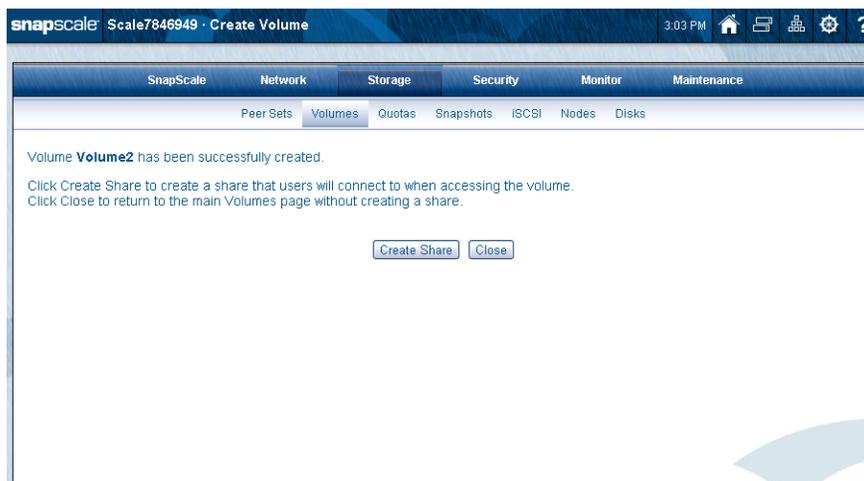


1. At **Storage > Volumes**, click the **Create Volume** button to open the **Create Volume** page.

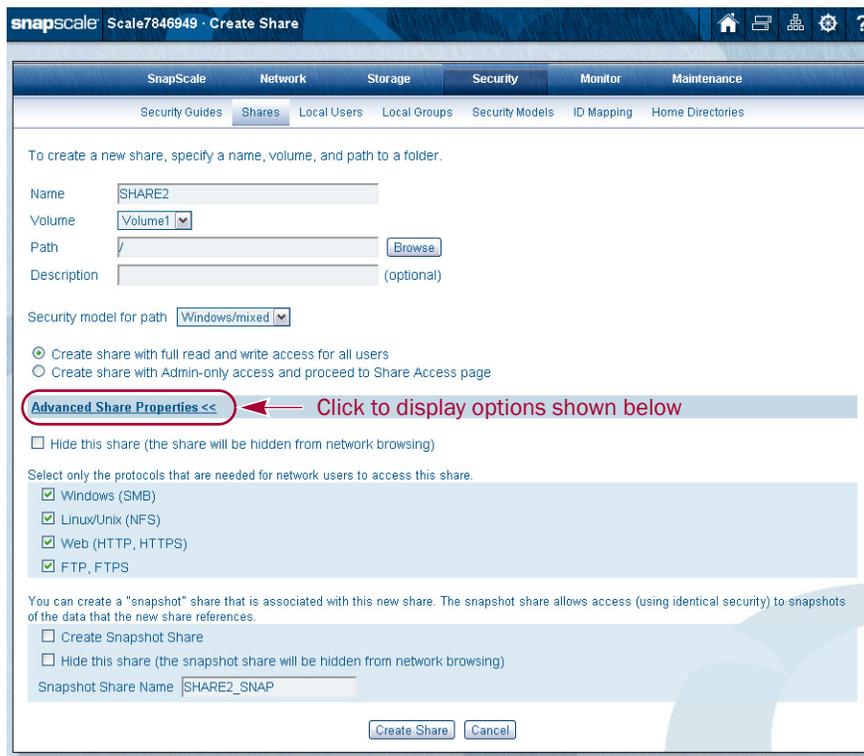


2. Make any necessary changes to the **options**.
 - It is recommended to enter a **Volume Name** to easily identify the specific volume.
 - If desired, keep the default of **No Limit** to allow the volume to consume an unlimited amount of cluster storage.
 - Otherwise, enter a **size**, changing the measurement units if needed.

- Click the **Create Volume** button again. A confirmation page is shown:



- At the confirmation page, click **Create Share** to create a share pointing to this volume (takes you to the **Shares** page).

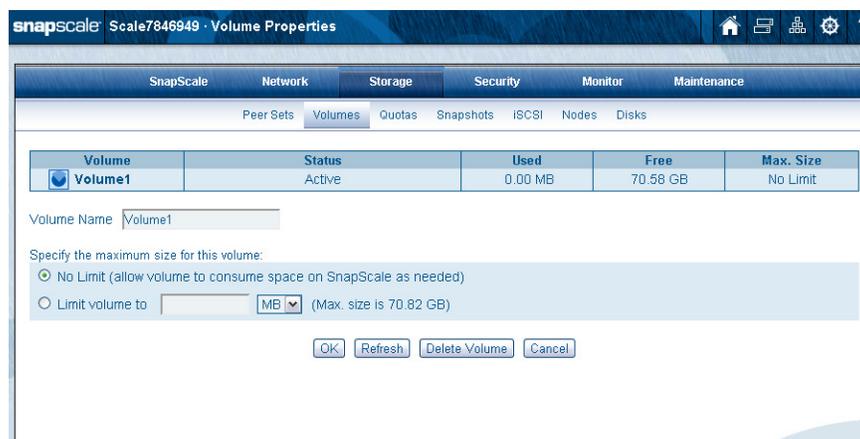


NOTE: The snapshot options at the bottom are only shown if snapshot space has been reserved.

- Enter the appropriate **data** and select the necessary options, then click **Create Share**. Additional options can be accessed by clicking the **Advanced Share Properties** link at the bottom. See [Shares](#) in [Chapter 6](#) for complete details.
- Click **Close** twice to return to the **Admin Home** page.

Edit Volume Properties

By clicking a volume's name on the main page, details of that particular volume are shown on the **Volume Properties** page.



From this secondary page, you can:

- Change the volume name.
- Set maximum volume size (specific limit or no limit).
- Delete the entire volume.

Rename a Volume

In the **Volume Name** field, enter a unique volume name of 32 alphanumeric characters and spaces, then click **OK**.

Specify Maximum Volume Size

There are two options controlling the maximum size of a volume:

- **No Limit** – This is the recommended option because it allows the volume to consume space as needed.
- **Limit Volume to** – Establish a maximum volume size limit by entering the amount and selecting a unit of measure (MB, GB, TB, or PB). The volume then grows in size until it reaches its maximum. If email notification has been enabled, alerts are sent as the maximum is approached. (To enable email notification, see [Email Notification in Chapter 8](#))

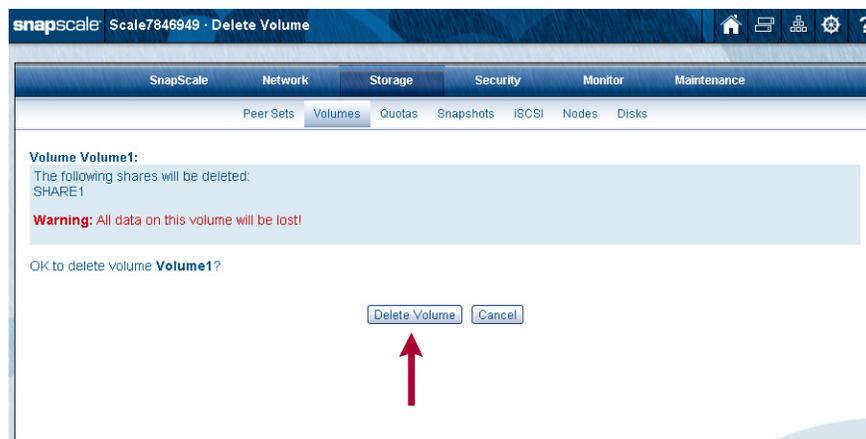
NOTE: If you reset the maximum size of a volume to less than its current size, the volume is treated as full and no more data can be written to it until the actual space consumed drops below the maximum size again. When done, click **OK**.

Deleting Volumes

To delete a volume, go to the **Volume Properties** page and click the **Delete Volume** button. At the confirmation page, click the **Delete Volume** button again. You are returned to the **Volumes** page and the volume is deleted in the background.



CAUTION: Deleting a volume deletes all the shares and data on the volume.



Quotas

Quotas are configured by accessing the **Storage > Quotas** page of the Web Management Interface. This default page shows all volumes on the cluster and their space/file quotas.



Assigning quotas ensures that no one user or group consumes a disproportionate amount of volume capacity measured by either space consumed or number of files created. The **Quotas** page also displays space consumed and files created by each user or NIS group regardless of whether a quota is applied to them, allowing for precise tracking of usage patterns. You can set individual quotas for any NIS, Windows domain, or local user known to the SnapScale. Group quotas are available only for NIS groups.

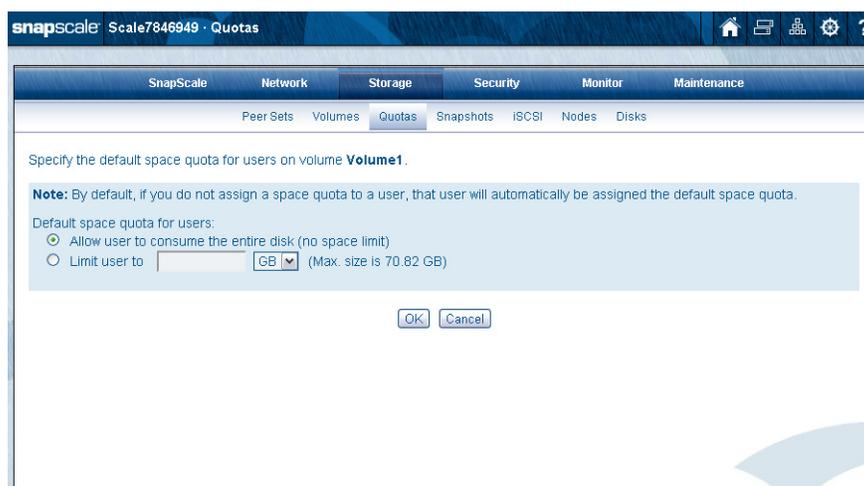
For users and groups, there are no pre-assigned default quotas on the SnapScale. When quotas are assigned, you can assign a default space or file quota for all users, or allow all users to have unlimited space on the volume. Unless you assign individual user or group quotas, all users and groups will receive the default quota.

In calculating usage, the SnapScale looks at all the files on the server that are owned by a particular user and adds up the file sizes. Every file is owned by the user who created the file and by the primary group to which the user belongs. When files are copied to the cluster, their size and count are applied against both the applicable user and NIS group quotas.

Default Quotas

On the main **Quotas** page (**Storage > Quotas**), the last two columns of the table show the default quotas for disk space and number of files. To change these settings, click the number (or no limit text) in the row under the default space or file quota column. A page is shown for the appropriate quota type options:

Default Space Quota Page



The screenshot shows the SnapScale web interface for the Quotas page. The breadcrumb trail is "Storage > Quotas". The page title is "Specify the default space quota for users on volume **Volume1**". A note states: "By default, if you do not assign a space quota to a user, that user will automatically be assigned the default space quota." Under "Default space quota for users:", there are two radio button options: "Allow user to consume the entire disk (no space limit)" (which is selected) and "Limit user to" followed by a text input field, a dropdown menu set to "GB", and the text "(Max. size is 70.82 GB)". At the bottom are "OK" and "Cancel" buttons.

To make changes, choose to either use the entire disk or a space of a specific size. For a specific size, enter the maximum amount and select the units. Click **OK** to accept.

Default File Quota Page

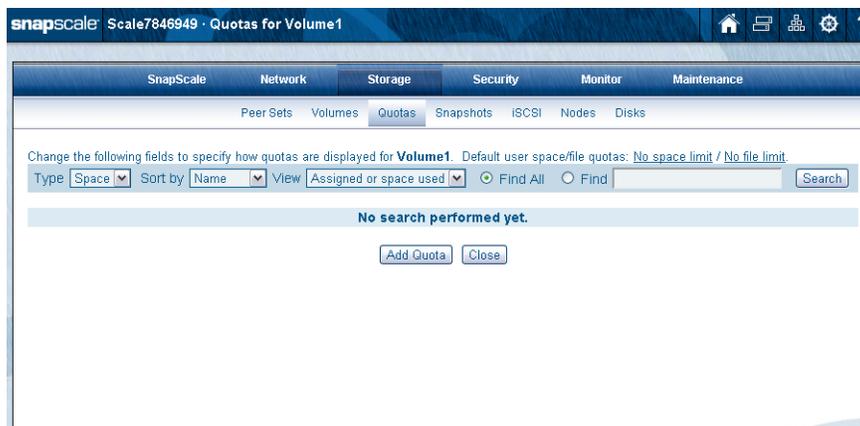


The screenshot shows the SnapScale web interface for the Quotas page. The breadcrumb trail is "Storage > Quotas". The page title is "Specify the default file quota for users on volume **Volume1**". A note states: "By default, if you do not assign a file quota to a user, that user will automatically be assigned the default file quota." Under "Default file quota for users:", there are two radio button options: "Allow user to create any number of files (no file limit)" (which is selected) and "Limit user to" followed by a text input field, the text "files.", and the text "(Max. assignable limit is 1,000,000,000 files)". At the bottom are "OK" and "Cancel" buttons.

To make changes, choose either to have no limit or a specific number of files. For a specific limit, enter the maximum number of files. Click **OK** to accept.

Quotas for Volume Page

From the **Quotas** page, you can create, view, or modify user and group quotas for a volume by clicking the volume's name in the **Volume** column on the far left. A **Quotas for Volume** page is displayed:



The page shows the available search and view options for the selected volume and the **Default user space/file quotas**. The two defaults shown can be either an amount or a text string:

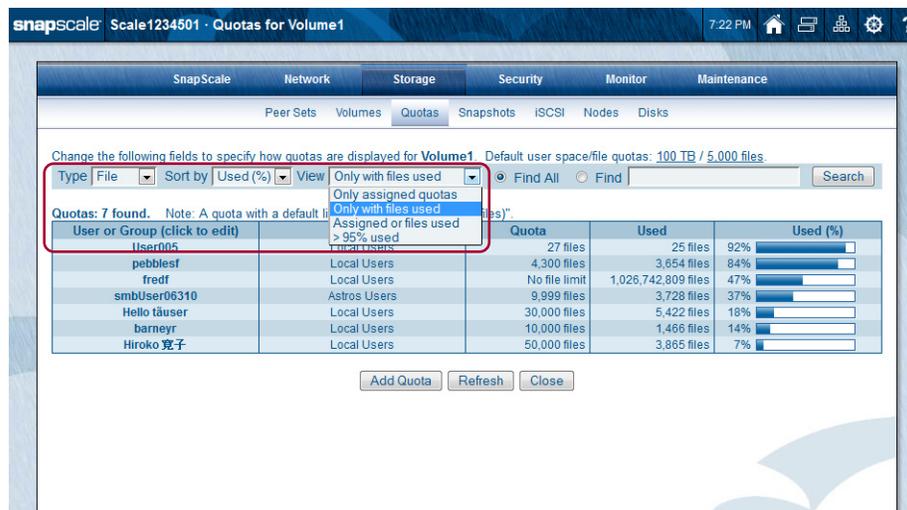
- **An amount** – the default quota size or file count assigned to users in that volume who do not have a specific quota assigned to them.
- **A text string** – the text strings **No space limit** and/or **No file limit** are displayed when quotas are enabled but the default space size and/or file count of no limits are configured for users in that volume. This means the users can consume the entire disk, or as many files as desired, respectively.

The space and file count limits also double as links to access the Web Management Interface pages where default space and file quotas can be configured.

Search for Quotas or Space Consumed by a User or NIS Group

To narrow down the list shown on this page or find a specific user or NIS group, first use the search bar just under **Default user space/file quotas**.

1. Select the **Type**, **Sort By**, and **View** parameters.



- **Type** – Choose **Space** or **File**.
 - **Sort by** – Select **Name**, **Limit**, **Used**, or **Used (%)**.
 - **View** – Choose one of these view options:
 - **Only assigned quotas**
 - **Only with files used / Only with space used** (depends on Type setting)
 - **Assigned or files used**
 - **> 95% used**
2. Select **Find All** or **Find**.
- When entering a search string for **Find**:
- Returned results will include all users and groups whose name **contains** the string entered.
 - To search a specific Windows or NIS domain, enter the domain name followed by a slash (/) or backslash (\) before the search string.
 - To search only local users and groups, enter “local” followed by a backslash (\) before the search string.
3. Click **Search**.
- A detailed list of users or NIS groups that match the parameters is displayed including the quota and space used numbers:

The screenshot shows the SnapScale web interface for managing quotas on Volume1. The top navigation bar includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The 'Storage' section is active, showing options for Peer Sets, Volumes, Quotas, Snapshots, iSCSI, Nodes, and Disks. Below this, there are controls to change display fields for Volume1, with a search bar and 'Find All' selected. A table displays the search results:

User or Group (click to edit)	Domain	Quota	Used	Used (%)
Albert	Local Users	2.00 GB	0.00 MB	0%

Buttons for 'Add Quota', 'Refresh', and 'Close' are located below the table.

NOTE: The search results returned may be automatically limited. Fine tune your search by using a more specific string to return a shorter list or the name desired.

Parentheses around a quota amount indicates the volume default quota is being used. If the volume's default quota is set to “no limit,” then “(no limit)” is displayed. If the volume's default quota is set to an actual value, such as 500GB, then “(500 GB)” is displayed.

No parentheses around the quota amount indicates a specific quota has been assigned. If the default quota limit is set to “no limit” but a particular user's or group's quota is set to 750GB, then “750 GB” is shown instead of the default “(no limit).”

The one exception to this is NIS groups. They don't use a volume default quota, so “no limit” (without parentheses) is shown.

4. To make **changes**, click the user or NIS group name.
See [Editing or Removing Quotas](#) on page 5-20.

Add Quota Wizard

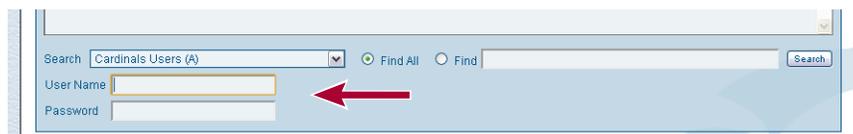
1. Click the **Volume name** link on the **Quotas** initial page to open the **Quotas for Volume** page for that volume.

- Click **Add Quota** to launch the search wizard.



- To search for a user or NIS group, select the local or domain option from the Search drop-down list, enter the search string (or select Find All), and click **Search**.

NOTE: For domains that require authentication (by showing an “(A)” after the name), after selecting the domain name, enter the User Name and Password for that domain.



- Returned results will include all users and NIS groups whose name **begins** with the string entered in the Search field.
- The search results returned may be limited. Fine tune your search by using a more specific string to return the names desired.
- On the rare occasion you need to search for a Windows domain that's not listed (“remote domain”), select a Windows domain from the Search drop-down list through which to search, then enter in the Find box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, `remote_domain\user_name`). After you click Search, another authentication prompt may be presented to authenticate with the remote domain.

- From the search results, select the appropriate **user or NIS group**, and click **Next** to show the configuration page for that user or group.

Quota for user **Local User/Al** on volume **Volume1** (Volume1 maximum size is 70.82 GB).

Quota	Used	Used (%)
(10 GB) (No file limit)	0.00 MB 0 files	0% 0%

User space quota:

- No space limit (user can consume space up to the maximum size of the volume)
- Limit to GB (Max. size is 70.82 GB)
- Use default user space quota (10 GB)

User file quota:

- No file limit (user can create any number of files)
- Limit to files. (Max. assignable limit is 1,000,000,000 files)
- Use default user file quota (No file limit)

OK Cancel

- Select or enter the desired **space or file quota** amounts, and click **OK**.

NOTE: NIS groups do not display the third option for using the default user space or file quota.

Editing or Removing Quotas

NOTE: Any changes override the default volume quota for this user or NIS group.

To edit or remove quotas of users or groups that have used space on this volume or have had specific quotas assigned to them from the volume:

- Click the **Volume name** link on the **Quotas** initial page to open the **Quotas for Volume** page for that volume.
- If necessary, **search** for a specific user to narrow the list to a more reasonable number. See [Search for Quotas or Space Consumed by a User or NIS Group](#) on page 5-17.

- To edit or remove the quota, from the search results select the appropriate **user or NIS group** from the left column to open the Quotas settings page.

Quota for user **Local User/Al** on volume **Volume1** (Volume1 maximum size is 70.82 GB).

Quota	Used	Used (%)
2 GB (No file limit)	0.00 MB 0 files	0% 0%

User space quota:

- No space limit (user can consume space up to the maximum size of the volume)
- Limit to (Max. size is 70.82 GB)
- Use default user space quota (10 GB)

User file quota:

- No file limit (user can create any number of files)
- Limit to files. (Max. assignable limit is 1,000,000,000 files)
- Use default user file quota (No file limit)

OK Cancel

- Select or enter the **quota** desired:
 - When **editing**, choose a limit or the default quotas.
 - To **remove** a specific quota limit, set both the space and file quotas to no limit.

NOTE: NIS groups do not display the third option for the default space or file quotas.

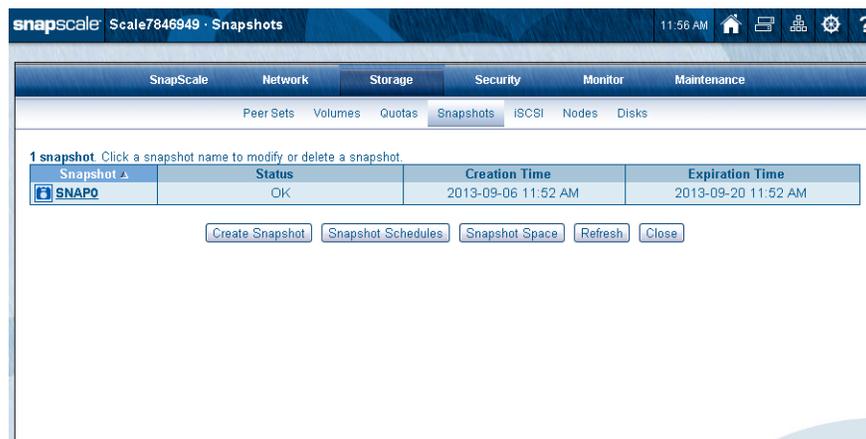
- Click **OK**.

Snapshots

A *snapshot* is a consistent, stable, point-in-time image of the cluster storage space that can be backed up independent of activity on the cluster storage. Snapshots can also satisfy short-term backup situations such as recovering a file deleted in error without resorting to tape. Perhaps more importantly, snapshots can be incorporated as a central component of your backup strategy to ensure that all data in every backup operation is internally consistent and that no data is overlooked or skipped.

NOTE: To preserve your cluster configuration and protect your data from loss or corruption, it is critical to schedule backups and snapshots.

Navigate to **Storage > Snapshots** in the browser-based Web Management Interface to create or schedule snapshots:



Snapshots Overview

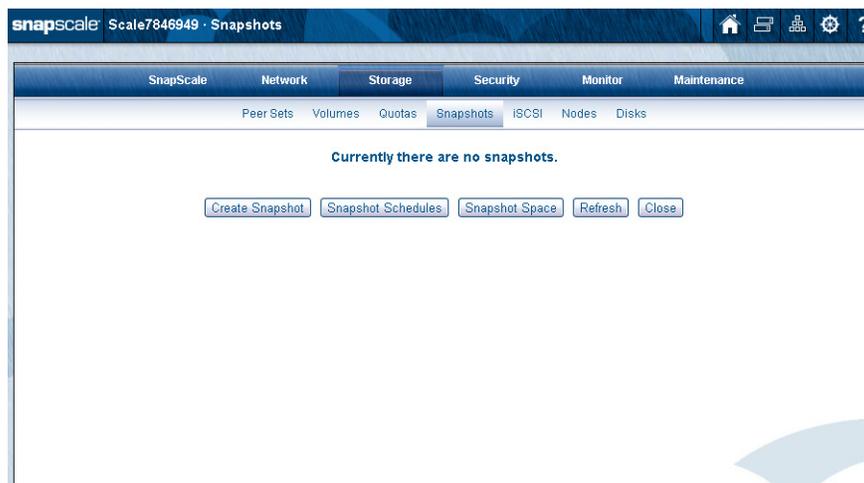
When working with snapshots, consider the following caveats:

- It is recommended that snapshots be taken when the system is idle or under low data traffic to minimize conflicts.
- Snapshots for the cluster storage space use snapshot space reserved on each peer set member drive. If no space is reserved (by unchecking the option box), snapshots are permanently disabled on the cluster.
- While 1% to 90% of the space can be reserved for snapshots, it is recommended that snapshot space be set to 20% of the cluster storage space during setup. Once set, the snapshot space can only be reduced, never increased.
- Snapshot space reserved from each peer set member drive is not necessarily identical to snapshot space of other drives in the same peer set. (This is most likely to occur if two or more drives in the same peer set have recently failed, even if they've been replaced with spares.) As a result, failure of a drive with unique snapshot data may cause one or more snapshots to be automatically deleted.
- Addition of a peer set to the cluster (including automatic peer set creation using new drives inserted into nodes or the addition of new nodes to the cluster) deletes all existing snapshots.
- Failure of a peer set deletes all snapshots.

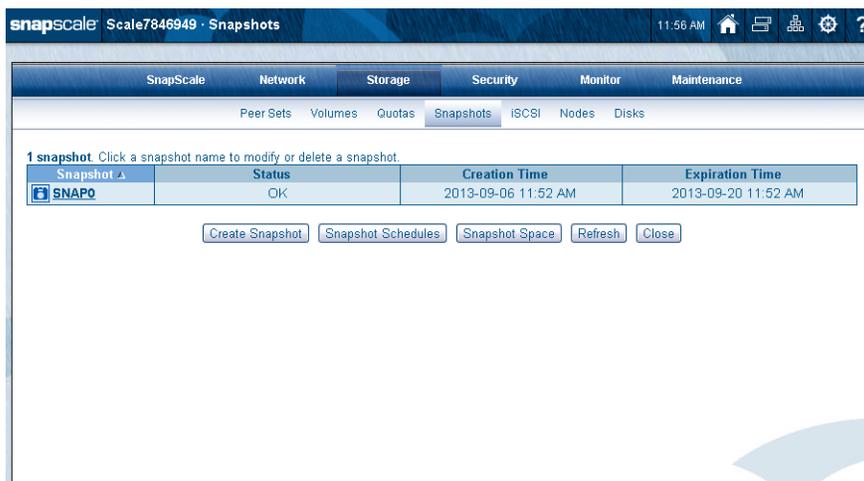
Creating Snapshots

Creating a snapshot involves naming, scheduling, and setting the duration of the snapshot. For regular data backup purposes, create a recurring snapshot. A recurring snapshot schedule works like a log file rotation, where a certain number of recent snapshots are automatically generated and retained as long as possible, after which the oldest snapshot is discarded. You can also create individual, one-time-only snapshots as needed.

If no snapshots are currently configured, you only see an empty page:



Once a snapshot is created, the page is populated with options for managing the snapshots:



These options are available in the Snapshots section of the Web Management Interface:

Action	Procedure
Create a New Snapshot	Click Create Snapshot . The process involves first defining snapshot parameters, and then scheduling when and how often to run the snapshot. Do not take more snapshots than your system can store, or more than 250 snapshots. Under normal circumstances, nine or ten snapshots are sufficient to safely back up any system.
Edit a Snapshot Schedule	Click the Snapshot Schedules button, and then click the snapshot name. You can modify all snapshot parameters.
Adjusting Snapshot Space Size	Specify the percentage of your SnapScale storage space to reserve for snapshots. NOTE: The storage space reserved for snapshots can be reduced, but it can never be increased once it is created.

Action	Procedure
Edit and Delete	Click the snapshot's name in the Snapshot column to open the Snapshot Properties page. You can edit the snapshot's name and duration, or delete the snapshot.
Refresh the Page	Clicking the Refresh button updates the page. This is helpful when waiting for a snapshot to complete.

When single snapshots are originally created or while recurring snapshots are active, the Refresh icon (🔄) is displayed on the right of the tab bar. It indicates that the snapshot data in the table is being refreshed every 5 minutes and can be clicked to manually refresh the data.

Clicking the **Close** button returns you to the **Storage Settings** page.

NOTE: The presence of one or more snapshots on a cluster can impact write performance. Additional snapshots do not have additional impact; in other words, the write performance impact of one snapshot on a cluster is the same as the impact of 100 snapshots.

Snapshots and Backup Optimization

When you back up a live volume directly, files that reference other files in the system may become out-of sync in relation to each other. The more data you have to back up, the more time is required for the backup operation, and the more likely these events are to occur. By backing up the snapshot rather than the volume itself, you greatly reduce the risk of archiving inconsistent data.

To Create a Snapshot

Using the Snapshots page in the Web Management Interface, you can create a snapshot now, later, or on a recurring schedule. When you select the Create Snapshot Later option, additional options are displayed.

Follow these steps to create a snapshot:

1. Go to **Storage > Snapshots**, and click **Create Snapshot**.
2. Enter or select the **options** for the snapshot:
 - a. Type in the **Snapshot Name** (20 character maximum).

- b. Specify **when** to create the snapshot.
 - Click **Create Snapshot Now** to run the snapshot immediately.
 - Click **Create Snapshot Later** to schedule the snapshot for a later time.

When you select the **Create Snapshot Later** button, a new input section appears below the option. Enter the Start Date and Start Time. Select either to create the snapshot only once (**One Time**) or to have it recurring periodically (**Recurring**) using an interval in hours, days, weeks, or months.

- c. Specify the **duration** of the snapshot.

NOTE: In the Duration field, specify how long the snapshot is to be active in hours, days, weeks, or months. The SnapScale automatically deletes the snapshot after this period expires, as long as no older unexpired snapshots exist on which it depends. If any such snapshot exists, its termination date is displayed at the bottom of the page. You must set the duration to a date and time after the displayed date.

3. Create the snapshot by clicking **Create Snapshot**.

If you elected to run the snapshot immediately, it appears in the Current Snapshots table. If you scheduled the snapshot to run at a later time, it appears in the Scheduled Snapshots table.

Adjusting Snapshot Space

NOTE: Once the SnapScale cluster is created, the storage space reserved for snapshots can only be decreased. It can never be increased.

If you have reserved storage space for snapshots during the setup of your cluster, you can use the **Snapshot Space** button to access the page where you can decrease the size of the space.



1. Go to **Storage > Snapshots**, and click **Snapshot Space**.
2. Reduce or remove the **reserved space**:
 - Using the drop-down list, choose a **lower percentage** of reserved space.
 - Uncheck the **reserve space for snapshots** box to release all reserved space.

CAUTION: Unchecking the reserve space box causes all the reserved space to be released, deletes all existing snapshots, and permanently disables snapshots on the cluster.

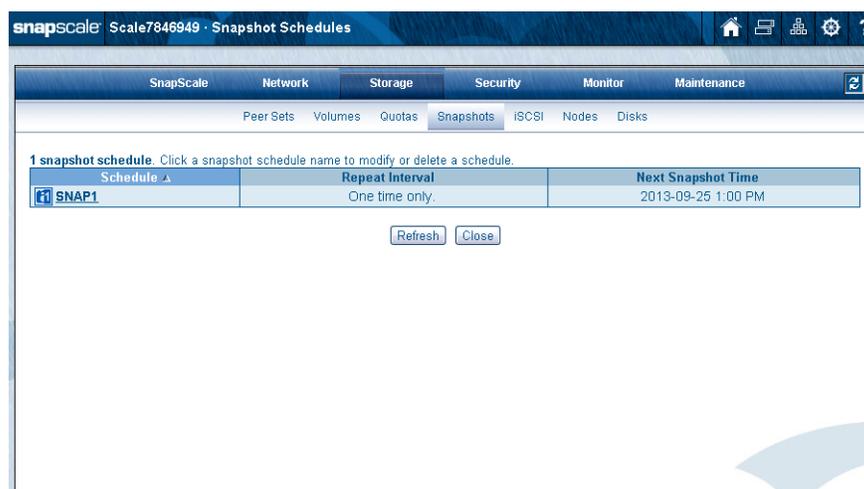
3. Click **OK** to complete the process.

Accessing Snapshots

After snapshots are created, they can be accessed via a snapshot share. Just as a share provides access to a portion of a live volume, a snapshot share provides access to the same portion of the filesystem on all current snapshots of the volume. The snapshot share's path into snapshots mimics the original share's path into the live volume. The snapshot share is created in the **Shares** section under the **Security** tab. See [Shares](#) in [Chapter 6](#) for details.

Scheduling Snapshots

To view when snapshots are currently scheduled to occur, click **Snapshot Schedules**:

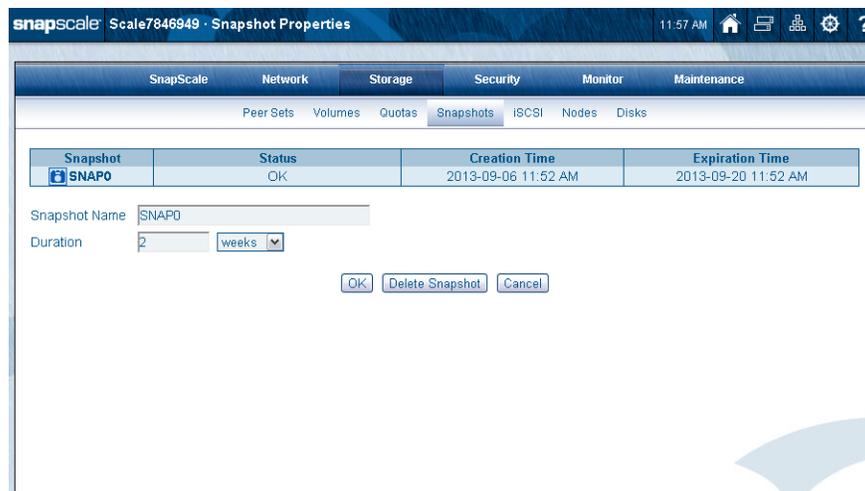


The **Snapshot Schedules** page shows a list of scheduled snapshots pending. **Repeat Interval** and **Next Snapshot Time** shows the details of when snapshots are scheduled to be taken.

Snapshots should ideally be taken when your system is idle. It is recommended that snapshots be taken before a backup is performed. For example, if your backup is scheduled at 4 a.m., schedule the snapshot to be taken at 2 a.m., thereby avoiding system activity and ensuring the snapshot is backed up.

Edit Snapshot Properties

From the **Snapshot** primary page table, you can click a snapshot name to access the **Snapshot Properties** page. There you can edit the name and duration, or delete the snapshot:



Edit a Snapshot

You can edit the name and duration by changing the data in the detail fields, and then clicking **OK**.

Delete a Snapshot

Click **Delete Snapshot** and then click it again on the confirmation page. The snapshot is deleted.

iSCSI Disks

Internet SCSI (iSCSI) is a standard that defines the encapsulation of SCSI packets in Transmission Control Protocol (TCP) and their transmission via IP. On SnapScale clusters, an iSCSI disk consumes cluster storage space as a single large file, but appears to a client machine as a local SCSI drive. This storage virtualization frees the administrator from the physical limitations of direct-attached storage media and allows capacity to be expanded easily as needed. Unlike standard volumes, SnapScale cluster iSCSI disks can be formatted by the iSCSI client to accommodate different application requirements.

Configuring iSCSI Initiators

Overland Storage has qualified a number of software initiators, PCI cards, and drivers to interoperate with SnapScale clusters. Refer to the vendor's documentation to properly install and configure you initiator to connect to the SnapScale iSCSI disks.

iSCSI Configuration on the SnapScale

iSCSI disks are created on the **Storage > iSCSI** page of the Web Management Interface. Before setting up iSCSI disks on your SnapScale cluster, carefully review the following information.

Basic Components of an iSCSI Network

iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. A basic iSCSI network has two types of devices:

- iSCSI initiators, either software or hardware, resident on hosts (usually servers), that start communications by issuing commands.
- iSCSI targets, resident on storage devices, that respond to the initiators' requests for data.

The interaction between the initiator and target mandates a server-client model where the initiator and the target communicate with each other using the SCSI command and data set encapsulated over TCP/IP.

Back up an iSCSI Disk from the Client, not the SnapScale

An iSCSI disk is not accessible from a share and thus cannot be backed up from the SnapScale cluster. The disk can, however, be backed up from the client machine from which the iSCSI disk is managed.

NOTE: While some third-party, agent-based backup packages could *technically* back up an iSCSI disk on the SnapScale cluster, the result would be inconsistent or corrupted backup data if any clients are connected during the operation. Only the client can maintain the filesystem embedded on the iSCSI disk in the consistent state that is required for data integrity.

iSCSI Multi-Initiator Support

Check the **Support Multiple Initiators** box to allow two or more initiators to simultaneously access a single iSCSI target. Multiple initiator support is designed for use with applications or environments in which clients coordinate with one another to properly write and store data on the target disk. Data corruption becomes possible when multiple initiators write to the same disk in an uncontrolled fashion.

NOTE: RAINcloudOS supports Windows 2003 and Windows 2008 Server failover clustering.

When the box for **Support Multiple Initiators** is checked, a warning message appears:

```
Uncontrolled simultaneous access of multiple initiators to the same
iSCSI target can result in data corruption. Only enable Multi-
Initiator Support if your environment or application supports it.
```

It functions as a reminder that data corruption is possible if this option is used when creating an iSCSI disk.

Disconnect iSCSI Disk Initiators before Shutting Down the Cluster

Shutting down the cluster while a client initiator is connected to an iSCSI disk appears to the client initiator software as a disk failure and may result in data loss or corruption. Make sure any initiators connected to iSCSI disks are disconnected before shutting down the cluster nodes.

iSCSI Disk Naming Conventions

iSCSI disks are assigned formal iSCSI Qualified Names (IQNs). These are used when connecting an initiator to an iSCSI target, and differ from the **iSCSI Disk Name** (alias) assigned when the iSCSI disk is created in the Web Management Interface. The full IQN is displayed for each iSCSI disk.

Click an iSCSI disk name to edit or delete the iSCSI disk. (Note: Mouseover a disk name to view its IQN.)

iSCSI Disk	Status ▲	Active Clients	IP Address	Authentication	Size
iscsi0	OK	0	10.25.11.163	None	10.00 GB
iqn.1997-10.com.snapscale:scale7846949:iscsi0		0	10.25.11.162	None	7.19 GB

The format of IQNs for new SnapScale iSCSI disks is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice]
```

where `[clustername]` is the name of the SnapScale cluster, and `[blockdevice]` is the internal identifier of the iSCSI disk on the target SnapScale cluster. All the `[blockdevice]` names are automatically created using the term `snapbd` appended with a sequence number (such as, `snapbd0`, `snapbd1`, and so on).

```
iqn.1997-10.com.snapscale:Scale1234567:snapbd0
```

Click an iSCSI disk name to edit or delete the iSCSI disk. (Note: Mouseover a disk name to view its IQN.)

iSCSI Disk	Status ▲	Active Clients	IP Address	Authentication	Size
iscsi0	OK	0	10.25.11.163	None	10.00 GB
iscsi1	OK	0	10.25.11.162	None	7.19 GB

The format of IQNs for **VSS-based iSCSI disks** on the SnapScale cluster is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice].[nnn]
```

where `[clustername]` is the name of the SnapScale cluster, `[blockdevice]` is the internal identifier of the iSCSI disk on the target SnapScale cluster, and `[nnn]` is a sequential number starting from 000. For example:

```
iqn.1997-10.com.snapscale:Scale1234567:snapbd0.000
```

The format of IQNs for **VDS-based iSCSI disks** on the SnapScale cluster is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice]
```

and the format for IQNs for **snapshots of iSCSI disks** on the SnapScale cluster is:

```
iqn.1997-10.com.snapscale:[clustername]:[blockdevice]-snap[n]
```

where, in both cases, `[clustername]` is the name of the SnapScale cluster, `[blockdevice]` is the internal identifier of the iSCSI disk on the target SnapScale cluster, and `[n]` is a sequential number starting from 0. For example:

```
iqn.1997-10.com.snapscale:Scale1234567:snapbd0-snap0
```

Create iSCSI Disks

Navigate to **Storage > iSCSI** and click **Create iSCSI Disk** to create, edit, or delete iSCSI disks on the SnapScale cluster. Be sure to read [iSCSI Configuration on the SnapScale](#) on page 5-27 before you begin creating iSCSI Disks.

The creation process involves first defining iSCSI parameters, then setting up security, and finally confirming your settings.

1. Navigate to **Storage > iSCSI** and click **Create iSCSI Disk**.
2. Enter the **iSCSI settings** for the disk name and size (16GB minimum).
Accept the default name or enter a new one. Use up to 20 alphanumeric, lowercase characters. Accept the default size of the remaining cluster space or enter a different size.
3. If you want your iSCSI Disk to allow **multiple** initiator connections, check that box.

NOTE: Data corruption is possible if this option is checked. See [iSCSI Multi-Initiator Support on page 5-28](#) for more information.

4. If desired, enable CHAP authentication by checking the **Enable CHAP Logon** box to display the hidden options.

Enter a **User Name** and **Target Secret** (password), and then confirm the password. Consider the following:

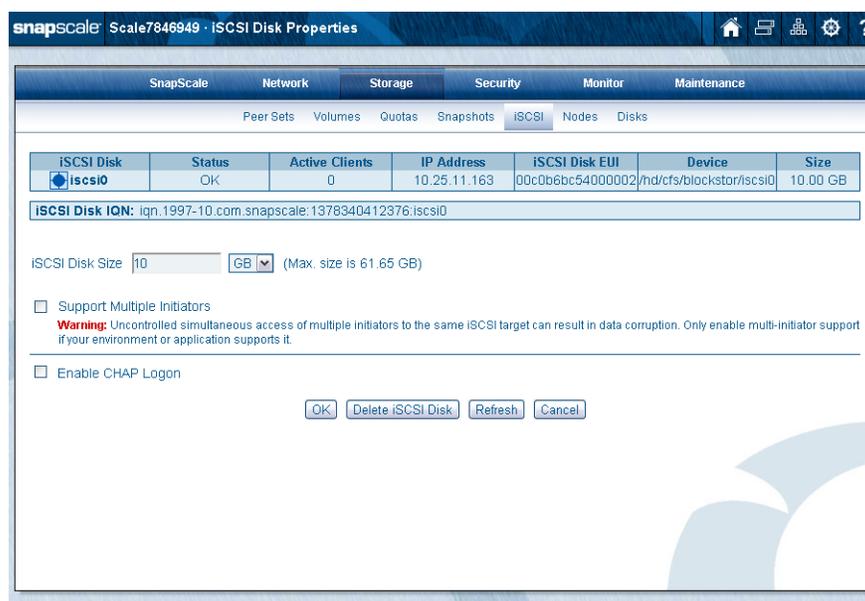
- Both items are case-sensitive.
 - The user name range is 1 to 223 alphanumeric characters.
 - The target secret must be a minimum of 12 and a maximum of 16 characters.
5. Click the **Create iSCSI Disk** button and, at the confirmation page, verify the settings.
 6. Click the **Create iSCSI Disk** button again to complete the process.
You are returned to the **iSCSI** page and the new iSCSI disk is displayed in the table there with the following information:

Label	Description
iSCSI Disk	The name of the iSCSI disk.
Status	Current condition of the iSCSI disk: <ul style="list-style-type: none"> • OK – The iSCSI disk is online and accessible. • Stopped – The iSCSI disk is currently stopped. • Failed – The iSCSI disk has failed.
Active Clients	The number of current sessions.
IP Address	The IP address used by the iSCSI disk.
Authentication	Either CHAP or None .
Size	The size of the iSCSI disk.

Edit iSCSI Disk Properties

NOTE: You cannot edit an iSCSI disk until all active clients have been disconnected from that disk. The hostname and IQN name of all connected initiators are displayed in the table.

After disconnecting all client initiators, click the iSCSI disk name in the table on the primary **iSCSI** page to display the **iSCSI Disk Properties** page.



1. On this **page**, you can do the following:
 - View the iSCSI Disk **IQN**.
 - Increase (but not decrease) the **size** of the iSCSI disk (if space is available).
 - Enable or disable support for **multiple initiators**.
 - Enable or disable **CHAP logon**.
2. Click **OK** to accept the changes (or **Cancel** to cancel).

CAUTION: The consistency of the internal filesystem on the iSCSI disk is primarily the responsibility of the file and operating systems on the iSCSI client used to format and manage the disk. Growing an iSCSI disk is handled differently by different operating systems and may lead to unexpected results on some client types.

Delete an iSCSI Disk

NOTE: You cannot delete an iSCSI disk until all active clients have been disconnected.

After disconnecting all client initiators, click the iSCSI disk name in the **iSCSI Disk** column to display the **iSCSI Disk Properties** page.



Click **Delete iSCSI Disk** (which is followed by a confirmation page) to delete the iSCSI disk.

Configuring VSS/VDS for iSCSI Disks

RAINcloudOS 4.0 provides VSS and VDS hardware providers to support Microsoft Volume Shadow Copy Services (VSS) and Virtual Disk Service (VDS) for iSCSI disks.

- The **VSS** hardware provider provides a mechanism for taking application-consistent Windows-native snapshots of iSCSI disks without performing full application (or system) shutdown. A snapshot of an iSCSI disk can be automatically created by a backup job run by a VSS-compatible backup application on a Windows initiator host, so that the job backs up the snapshot volume rather than the main production volume.

NOTE: VSS iSCSI snapshots are managed by the Windows client and represent the iSCSI disk, not the Snap volume on which the iSCSI disk resides. They are not related to RAINcloudOS snapshots as described in [Snapshots on page 5-21](#). The VSS iSCSI snapshot rollback feature is not currently supported.

- The **VDS** hardware provider allows administrators to natively manage SnapScale cluster iSCSI disks, using any VDS-compliant management console application.

Backing up an iSCSI Disk using VSS Snapshots

Windows VSS-compatible backup applications can create snapshots of SnapScale cluster iSCSI disks to perform consistent backups of application data without stopping the application, using the snapshot instead of the live volume as the backup source.

Each VSS snapshot of an iSCSI target consumes additional cluster storage space. The required space is 10% of the size of the iSCSI disk per snapshot. If this amount of free space is not available on the pool or volume, the VSS snapshot will not be created and an error will be reported by the SnapScale cluster VSS hardware provider to the Windows event log.

When creating iSCSI disks for later VSS snapshot use, be sure to leave at least 10% of the size of the iSCSI target free on the cluster.

NOTE: VSS snapshots can only be taken of Windows volumes that fully consume the iSCSI disk. Snapshots of iSCSI disks that contain multiple Windows volumes are not supported.

1. Add the **VSS client** to the SnapScale cluster.



- a. From the **Storage > iSCSI** page, click the **VSS/VDS Access** button.
- b. Click **Add**.
- c. Add the **hostname** of the VSS client you wish to grant access and click **Add** (the hostname is not case-sensitive).
The client hostname should appear in the VSS/VDS Clients box.

NOTE: Use the short hostname (*myclientname*) of the client only. Do not use the IP address or fully-qualified name (for example, *myclientname.mydomain.com*).

- d. When you have finished adding VSS clients, click **OK**.
2. Install the **VSS hardware** provider on the Windows iSCSI client.
 - a. Depending on the Windows client, locate *SnapServerToolsInstall32.exe* or *SnapServerToolsInstall64.exe* on the Overland website:
<http://docs.overlandstorage.com/snapserver>
 - b. Double-click the **executable** (.exe) to run the Installation Wizard on the VSS client and select the VSS/VDS hardware providers option. This will add the hardware provider to the Windows iSCSI client.
 3. Configure VSS-based **backups** of the iSCSI disk.
 - a. Connect the client **iSCSI initiator** to the Snap iSCSI disk and create a volume (if necessary). Add data or configure applications to use the iSCSI volume for the data repository.
 - b. Configure a VSS-based **backup** of the iSCSI disk. Where applicable, choose to use the SnapScale cluster VSS hardware provider in the backup job configuration. When the backup job is run, the snapshot of the iSCSI disk is automatically created and hosted by the SnapScale cluster as a virtual iSCSI disk (named after the main iSCSI disk with *snap[n]* appended), and the backup application performs the backup using the snapshot iSCSI disk. The snapshot will be deleted after the backup completes.

NOTE: VSS snapshots are not supported on SnapScale cluster iSCSI disks that have been configured into multiple Windows volumes.

Creating and Managing iSCSI LUNs Using VDS

1. Create the **volume** for the iSCSI disk on the SnapScale cluster using the Web Management Interface (**Storage > Volumes**).
The volume must be created on the SnapScale cluster before the iSCSI disk can be created using a VDS application such as Microsoft's *Storage Manager for SANs*.
2. Add **VDS clients** to the SnapScale cluster.
 - a. From the **Storage > iSCSI** page, click the **VSS/VDS Access** button.
 - b. Click **Add**.
 - c. Add the hostname of the VDS client you wish to grant access and click **Add** (the hostname is not case-sensitive). The client hostname should appear in the VSS/VDS Clients list.

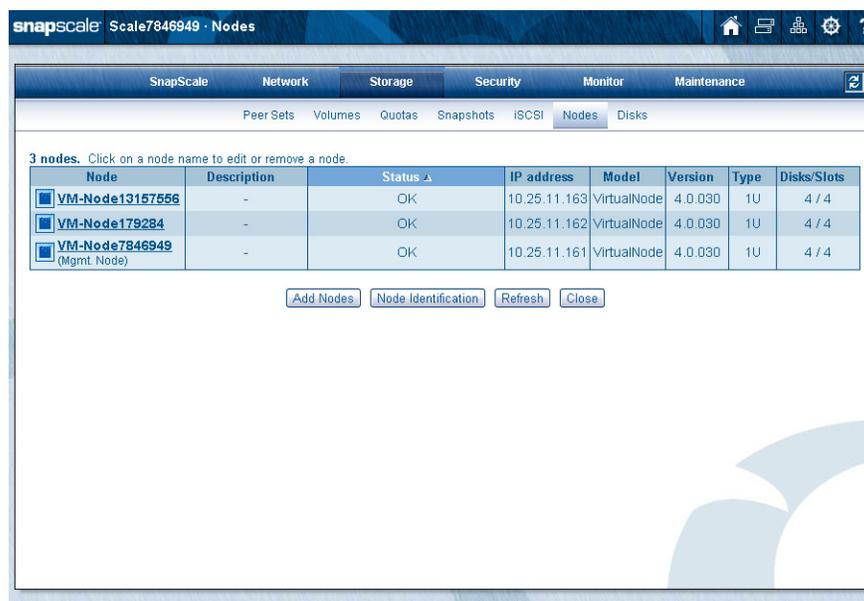
NOTE: Use the short hostname (*myclientname*) of the client only. Do not use the IP address or fully-qualified name (for example, *myclientname.mydomain.com*).
 - d. When you have finished adding VDS clients, click **OK**.
3. Install the **VDS hardware provider** on the Windows client.
 - a. Depending on the Windows client, locate *SnapServerToolsInstall32.exe* or *SnapServerToolsInstall64.exe* on the Overland website:
<http://docs.overlandstorage.com/snapservers>
 - b. Run the **Installation Wizard** on a VDS client and select the VSS/VDS hardware providers option. This will add the hardware provider to the Windows client.
4. Create and configure the **iSCSI disk** using *Storage Manager for SANs* (or other VDS-compliant application).

Deleting VSS/VDS Client Access

1. From the **Storage > iSCSI** page, click the **VSS/VDS Access** button.
2. Select the **VSS/VDS client** you want to delete from the VSS/VDS Clients list, and click **Delete**.
3. Click **OK** to save your changes.

Nodes

Use the **Nodes** page to manage the nodes that make up the cluster.



From this page, you can:

- Add a new node.
- Edit or delete the node (by clicking the node name to access the **Node Properties** page).
- Identify physical nodes via flashing LEDs.

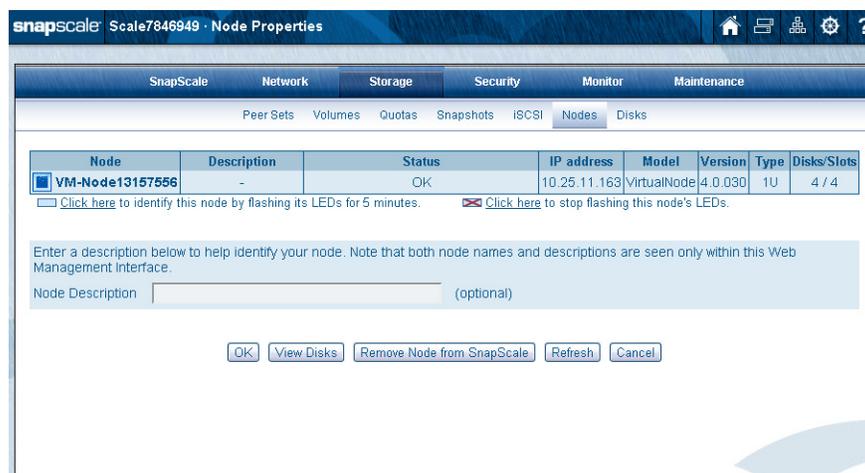
Nodes Overview

Some important points about SnapScale nodes:

- Users can access the cluster storage over any of the configured network protocols by connecting to any of the nodes.
- Because the storage space is unified across the cluster, connecting to any of the nodes provides access to the same data as any other cluster node.
- To balance network client access to the nodes, enter an **A** record to the DNS pointing to the cluster name for each IP address in the node IP address range. The DNS then uses round-robin name resolution requests for the cluster name among the node IP addresses. Alternatively, manually distribute clients accessing the cluster to different IP addresses in the node IP address range.
- When a node fails, the IP address it uses is automatically reassigned to another node. Clients connected to that IP address are forwarded to the new node, though this may cause a momentary interruption to storage access.
- Files opened by clients connected to any node are recognized by all nodes, and file locks are respected by all nodes.

Edit Node Properties

By clicking a node's name on the main **Nodes** page in the **Node** column of the table, details of that particular node are shown on a **Node Properties** page.



From this page, you can:

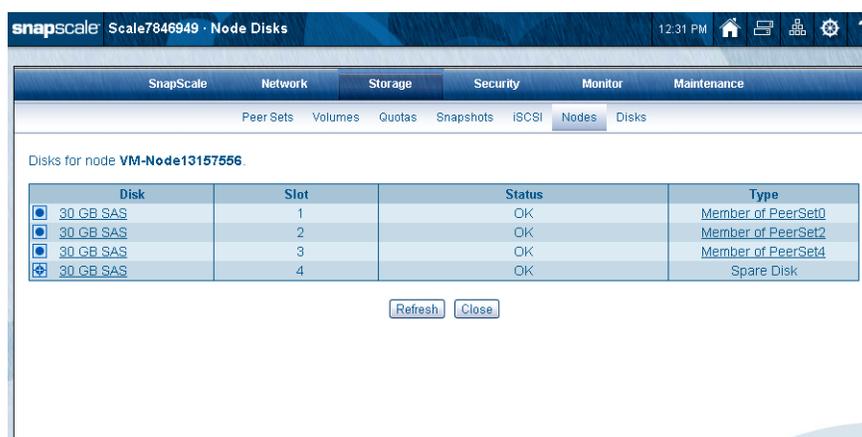
- Flash the node drive LEDs to help identify the node.
- Change the node description.
- View the drives in the node.
- Remove the node from the SnapScale cluster.

Flash the Node LEDs

Click the light-blue box () under the node name to start the LEDs flashing for up to five minutes. Click the box with the red "X" () to stop flashing the LEDs.

Node Drives

To view the drives that are installed in the node, click the **View Disks** button on the properties page.



When you click a drive's name in the **Disk** column on the **Node Disks** page, the **Disks** page is displayed with the physical location of the disk drive. See [Disks](#) on page 5-43 for more information.

Clicking the member name in the **Type** column takes you to the **Peer Sets** page with that member highlighted. See [Peer Sets Page](#) on [page 5-5](#) for more information.

When done, click **Close** to return to the **Node Properties** page.

Adding Nodes

A SnapScale cluster can also be expanded by adding more nodes. Expansion kits are available that consist of either two or three additional nodes and all the necessary cables. Documentation is included with each node that details how to install, cable, and power up the new node.

Once installed in a rack, clicking the **Add Node** button on the **Nodes** page starts a wizard to add one or more nodes to the cluster to expand the storage space. By default, all eligible nodes are pre-selected.

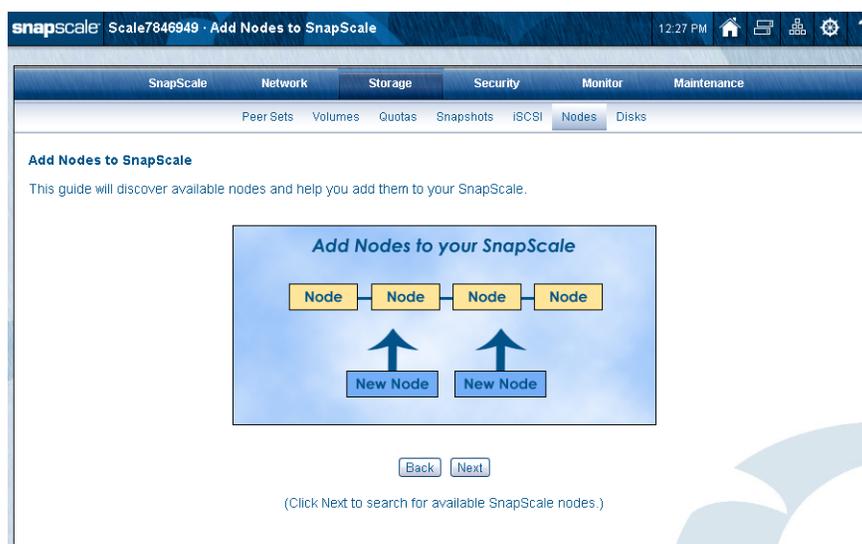
 **IMPORTANT:** In order to expand storage by adding nodes, the cluster must be able to create peer sets with each member on different nodes. To efficiently increase cluster storage, it is recommended that the number of new nodes you add is equal to or greater than the Data Replication Count (2x or 3x). It is highly recommended that these new nodes all be added in the same operation.

When adding nodes, all the following must be taken into consideration:

- The nodes must all be running the same version of RAINcloudOS (ROS). See [OS Update](#) in [Chapter 8](#).
- All the nodes, those already part of the cluster and those being added, must be attached to the same Client subnet.
- No expansion units can be attached to a node.
- The appropriate ports on the node must be available to create the proper bonding. (See [Client and Storage Networks](#) on [page 1-5](#).)

Follow these steps using the wizard to add your nodes:

1. Click **Add Nodes**.



2. Click **Next** to display node choices (Wizard Step 1):

Add Nodes to SnapScale: Select Nodes to Add

Step 1 | Step 2 | Step 3

Select the nodes below that you want to add to your SnapScale and click Next. (All eligible nodes are selected by default.)

Note: All nodes in the SnapScale must have identical RAINcloudOS (ROS) versions, and the client network interface for all nodes must be located on the same subnet.

Important: The data replication count is set at 2x. In order to prevent possible peer set synchronization issues you should select a minimum of 2 nodes to add to the SnapScale. Also, it is recommended that you select all nodes to add now, rather than adding nodes incrementally.

2 Eligible Nodes: (Note: 4 nodes are not eligible to be added to this SnapScale.)

Node	Model	ROS Version	Disks	Add to SnapScale
<input checked="" type="checkbox"/> Node1234511	X2	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: (No Disk) 11: (No Disk) 12: (No Disk)	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Node1234512	X4	4.0.119	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: 1.95 TB 11: 3.91 TB 12: 3.91 TB 13: 1.95 TB 14: 1.95 TB 15: 3.91 TB 16: 3.91 TB 17: 1.95 TB 18: 1.95 TB 19: 3.91 TB 20: 3.91 TB 21: 1.95 TB 22: 1.95 TB 23: 3.91 TB 24: 3.91 TB Rear Disks 25: 1.95 TB 26: 1.95 TB 27: 3.91 TB 28: (No Disk) 29: (No Disk) 30: (No Disk) 31: (No Disk) 32: (No Disk) 33: (No Disk) 34: (No Disk) 35: (No Disk) 36: (No Disk)	<input checked="" type="checkbox"/>
<input type="checkbox"/> Node1234502	X4	4.0.077	1: 1.95 TB 2: 1.95 TB 3: 3.91 TB 4: 3.91 TB 5: 1.95 TB 6: 1.95 TB 7: 3.91 TB 8: 3.91 TB 9: 1.95 TB 10: 1.95 TB 11: 3.91 TB 12: 3.91 TB 13: 1.95 TB 14: 1.95 TB 15: 3.91 TB 16: 3.91 TB 17: 1.95 TB 18: 1.95 TB 19: 3.91 TB 20: 3.91 TB 21: 1.95 TB 22: 1.95 TB 23: 3.91 TB 24: 3.91 TB Rear Disks 25: 1.95 TB 26: 1.95 TB 27: 3.91 TB 28: (No Disk) 29: (No Disk) 30: (No Disk) 31: (No Disk) 32: (No Disk) 33: (No Disk) 34: (No Disk) 35: (No Disk) 36: (No Disk)	<input type="checkbox"/> (Different/older ROS version; click here to upgrade this node.)

Back | Re-Detect Available Nodes | Next

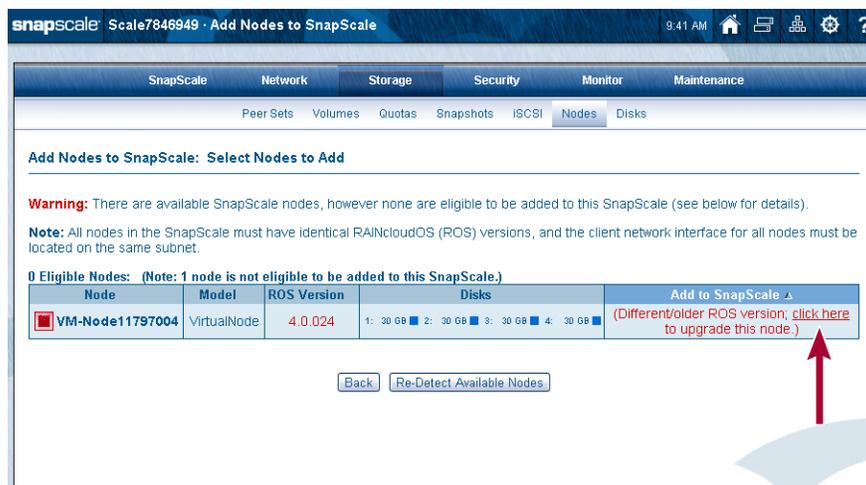
3. At Wizard Step 1, check the boxes for the **nodes** you want to add to the cluster, and click **Next**.

NOTE: By default, all eligible nodes are pre-selected. It is recommended to accept all the nodes to ensure the optimum configuration.

If the node bond type doesn't match the cluster bond type, special informational messages are shown in the wizard based on the existing situation:

- If one or more of the nodes have a different Storage network bond type than the cluster, a note table is displayed in Wizard Step 1 showing the bond issues (or other errors).
- The final wizard page displays a percent completed status while the updated nodes with changed bond types are being rebooted.

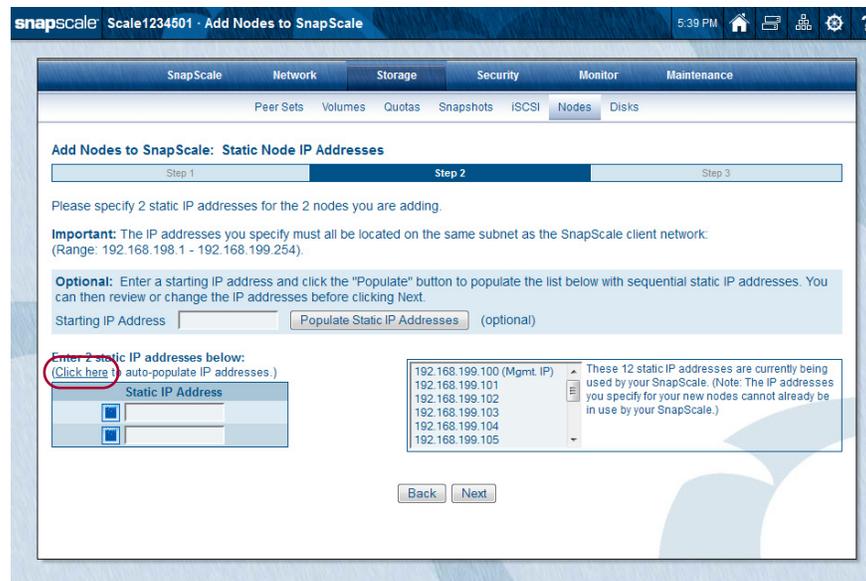
If an available node is not eligible to be added to the cluster due to a firmware mismatch, click the **click here** link, then login to the node and perform an update to the firmware (see [OS Update](#) in Chapter 8, "Maintenance"). When updated, go back to the browser tab showing the cluster and click **Re-Detect Available Nodes**.



- At Wizard Step 2, configure the **static IP addresses** for the nodes, and click **Next**. It is recommended that you click the **Click here** option at the lower left to automatically add IP addresses based on the addresses being currently used.

NOTE: When more new nodes are being added to the cluster than there are unused IP addresses in the node address pool, more IP addresses must be added to the pool.

You can also enter a starting address in the Populate field based on the static IP addresses (in the list on the right) currently being used by your SnapScale cluster, and then click the **Populate Static IP Addresses** button.



5. At Wizard Step 3, verify the data and click **Add Nodes to SnapScale**.

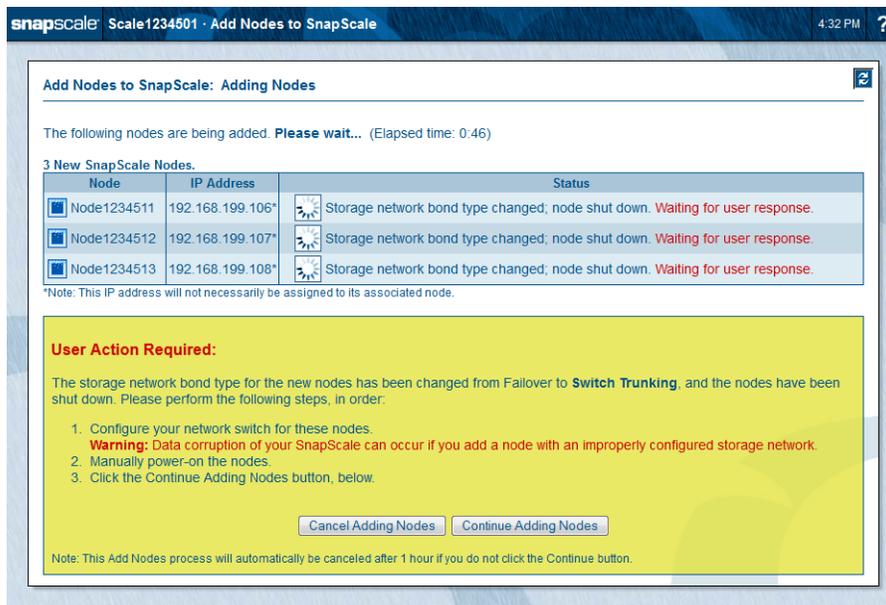


6. At the confirmation page, click **Add Nodes to SnapScale** again.



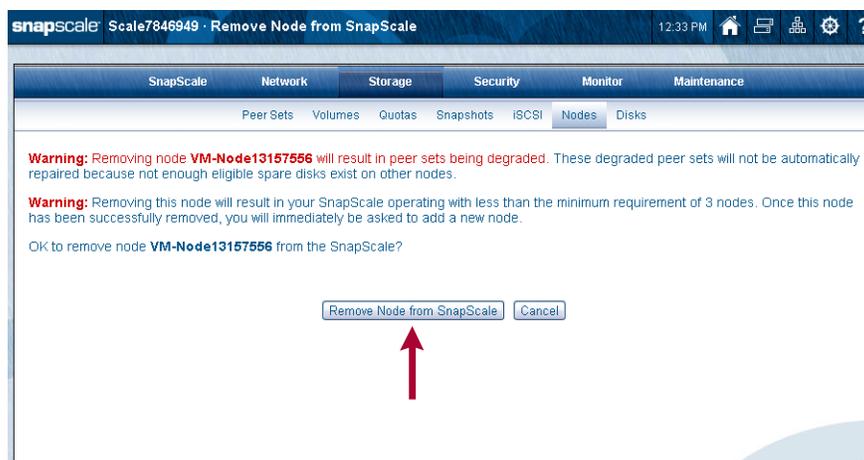
The new nodes are added to the cluster and the peer sets are built. This process takes several minutes.

IMPORTANT: If the Storage network bond type on the new nodes needs to be changed to or from Switch Trunking or Link Aggregation (802.3ad) to match the cluster bond type, an additional step is shown after shutting down the node so you can reconfigure the network switch and restart the node. You can also cancel the adding of the nodes if desired. However, if no action is taken in the first hour, after that time the adding of the nodes is automatically canceled.



Removing Nodes

To remove a node from a SnapScale cluster, go to the **Node Properties** page and click the **Remove Node from SnapScale** button. At the confirmation page, click the **Remove Node from SnapScale** button again.



IMPORTANT: Removing a node may result in one or more peer sets becoming degraded. These degraded peer sets may not be automatically repaired if there are not enough eligible spare drives on other nodes. Removing this node may also result in your SnapScale cluster operating with less than the minimum requirement of 3 nodes.

NOTE: If removing the node would destroy one or more peer sets, an error message is returned and the node is not removed.

The node itself is no longer associated with the cluster and becomes an Uninitialized node that can be added to another cluster.

Node Identification

The **Node Identification** page (accessed by the button on the main **Nodes** page) provides a convenient place to check the nodes and optionally change their descriptions for easier identification in the Web Management Interface.

NOTE: These options are also accessible from the Node Properties page.

Click a light-blue box () next to the node name to start the node's LEDs flashing for up to five minutes. Click the box with the red "X" () to stop flashing the LEDs or click the link next to the same icon below the nodes table to stop all node LEDs flashing.

3 nodes. Click this icon to identify a node by flashing its LEDs for 5 minutes. Click this icon to stop flashing a node's LEDs.

Node	Description	IP address	Model	Type	Disks/Slots
<input type="checkbox"/> VM-Node13157556 <input checked="" type="checkbox"/>		10.25.11.162	VirtualNode	1U	4 / 4
<input type="checkbox"/> VM-Node179284 <input checked="" type="checkbox"/>		10.25.11.163	VirtualNode	1U	4 / 4
<input type="checkbox"/> VM-Node7846949 <input checked="" type="checkbox"/>		10.25.11.161	VirtualNode	1U	4 / 4

[Click here to stop flashing LEDs on all nodes.](#)

OK Cancel

Stop All

Disks

The **Disks** page is a graphic representation of the peer sets and disk drive status on your cluster. The legend on the page explains the meaning of each icon.

Move the mouse over a disk icon to highlight all disks in the disk's peer (or spare) set. Click a disk icon to view disk details. Click to flash (for approx. 5 minutes) a node's LEDs for identification. Click to stop flashing LEDs.

Members of the Same Peer Set

Legend: Disk OK Disk Unused Disk Failure Empty Slot (No Disk) Spare Disk Spare Too Small for Some Peer Sets Spare Too Small for Any Peer Set

[Click here to stop flashing LEDs on all nodes.](#) **Stop All**

- Click a drive icon () to view drive details.
- Hover over a drive icon () to highlight in blue all the other members of the peer set.
- Hover over a spare drive icon () to view other spare drives.
- Click a unit's LED icon () to flash the unit's status and drive status LEDs for identification. The LEDs flash amber. Click the LED stop icon () to stop the unit's LEDs from flashing.

NOTE: The LEDs continue to flash for five minutes unless stopped. To stop flashing LEDs for all units, click the link next to the stop icon located below the Legend list.

Replacing Drives

Should a drive fail (solid red LED), it can be replaced (hot-swapped) without shutting down the SnapScale node. If a spare is available on a node that doesn't already have an active member of the failed drive's peer set, the spare automatically replaces the failed drive and the new drive being installed automatically becomes a spare. If no spares are available, the new drive automatically becomes a member of the failed drive's peer set.

A failed drive can be removed and replaced anytime. When hot-swapping multiple drives, it is recommended to swap one drive at a time to avoid timing issues.

NOTE: Hot-removed (non-spare) drives cannot be added directly back into the peer set from which they were removed. When any non-spare drive is physically removed, it is also removed from the peer set to which it belonged. That peer set then becomes degraded and it attempts to incorporate a suitable cluster spare. If the removed drive is reinserted and there is a degraded peer set or not enough spares to satisfy the spare count, the reinserted drive is reconfigured as a spare; otherwise, the reinserted drive becomes unused and available to the user to incorporate manually via the Web Management Interface.

If there are no errors, after the new drive is incorporated, any alert LEDs are turned off and system statuses are updated.

Adding Drives



CAUTION: If new peer sets are created when adding new drives to a node in a SnapScale cluster, all existing snapshots will be deleted.

NOTE: SnapScale only supports SAS (and nearline SAS) drives.

If empty slots are available on one or more nodes, you can add Overland-approved drives to either expand cluster storage space or add hot spares. When adding drives to expand storage space, distribute the new drives evenly across all the cluster nodes.

NOTE: Drives should be added without shutting down the node so that the cluster properly recognizes each drive. Note that drives with different rotational speeds cannot be combined in the same cluster.

In order to properly create peer sets with each member on different nodes, if you have the Data Replication Count set at 2x, you must add drives in groups of two, each one to different nodes. For a 3x count, add drives in groups of three, each one in a different nodes.

Once the new drives are added to the nodes, they must be incorporated using the **New Disks Detected** page in the Web Management Interface to enable the cluster to use them properly to create new peer sets and spares.

Important Considerations

When adding new drives, consider the following:

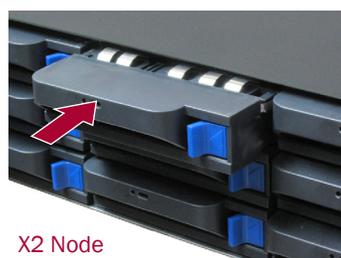
- Only add new drives (hot-insert) when the cluster is up and fully operational. Do not add new drives when the cluster is powered down.
- New drives can be added to Uninitialized nodes either while running or powered down.
- If there are fewer spares in the cluster than the spare count specifies, drives added to any cluster node are automatically configured as hot spares until the spare count is satisfied.
- If there is a degraded peer set on the cluster when adding a new drive and there are no existing spares, the drive will automatically be incorporated into the peer set as long as it is not on one of the nodes containing another active member of the peer set.
- When replacing a failed drive, it is recommended that the new drive be installed in the same slot as the old one to maintain a capacity balance.
- When adding a drive that replaces a failed drive in a peer set, the **Peer Sets** page will display that peer set as rebuilding the new drive into the peer set.

Drive Installation

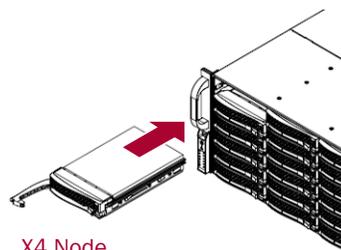
NOTE: Do not remove the disk drives from their carriers. Doing so voids the drive warranty. Unless adding drives to an Uninitialized nodes, the cluster must be up and fully operational.

Install the drives into the available slots:

1. Remove the **blank drive carriers** from the slots that will be used for the new drives (leaving the remaining blank carriers in place).
2. Positioning a **drive carrier** in front of the appropriate **bay**:
 - For the **X2** node, slide it in until the **latch** clicks, locking the assembly in the bay.
 - For the **X4** node, slide it in until it stops and then close the **latch** to lock it in place.



X2 Node



X4 Node

3. Repeat **Step 2** for **each** remaining drive carrier being installed.



IMPORTANT: To maintain proper airflow and cooling, a drive carrier or a blank carrier must be installed in every slot. No empty slots are allowed.

On the **Disks** page, any newly detected drives show a disk unused icon next to the drive. It may take a minute or two before the drives appear as unused and the new disks detected banner is displayed. The alert link in the banner takes you to the **New Disks Detected** page for incorporation.

snapScale Scale1234501 · Disks 5:45 PM

SnapScale Network Storage Security Monitor Maintenance

Peer Sets Volumes Quotas Snapshots iSCSI Nodes Disks

New disks detected. [Click to view and incorporate disks.](#)

Move the mouse over a disk icon to highlight all disks in the disk's peer (or spare) set. Click a disk icon to view disk details. Click to flash (for approx. 5 minutes) a node's LEDs for identification. Click to stop flashing LEDs.

Node	Slot	Capacity	Status
Node1234501 (X2, 16.8 TB)	1	0.98 TB SATA	Disk OK
	2	0.98 TB SATA	Disk OK
	3	1.95 TB SATA	Disk OK
	4	1.95 TB SATA	Disk OK
	5	0.98 TB SATA	Disk OK
	6	0.98 TB SATA	Disk OK
	7	1.95 TB SATA	Disk Unused
	8	1.95 TB SATA	Disk OK
	9	0.98 TB SATA	Disk OK
	10	(No Disk)	Empty Slot (No Disk)
	11	1.95 TB SATA	Disk OK
	12	1.95 TB SATA	Disk OK
Node1234503 (X2, 14.85 TB)	1	0.98 TB SATA	Disk OK
	2	0.98 TB SATA	Disk OK
	3	1.95 TB SATA	Disk OK
	4	1.95 TB SATA	Disk OK
	5	0.98 TB SATA	Disk OK
	6	0.98 TB SATA	Disk OK
	7	1.95 TB SATA	Disk OK
	8	1.95 TB SATA	Disk OK
	9	0.98 TB SATA	Disk OK
	10	(No Disk)	Empty Slot (No Disk)
	11	(No Disk)	Empty Slot (No Disk)
	12	2.15 TB SATA	Disk OK
Node1234507 (X2, 12.7 TB)	1	0.98 TB SATA	Disk OK
	2	0.98 TB SATA	Disk OK
	3	1.95 TB SATA	Disk OK
	4	1.95 TB SATA	Disk OK
	5	0.98 TB SATA	Disk OK
	6	0.98 TB SATA	Disk OK
	7	1.95 TB SATA	Disk OK
	8	1.95 TB SATA	Disk OK
	9	0.98 TB SATA	Disk OK
	10	(No Disk)	Empty Slot (No Disk)
	11	(No Disk)	Empty Slot (No Disk)
	12	(No Disk)	Empty Slot (No Disk)
Node1234509 (X2, 14.85 TB)	1	0.98 TB SATA	Disk OK
	2	0.98 TB SATA	Disk OK
	3	1.95 TB SATA	Disk OK
	4	1.95 TB SATA	Disk OK
	5	0.98 TB SATA	Disk OK
	6	0.98 TB SATA	Disk OK
	7	1.95 TB SATA	Disk OK
	8	1.95 TB SATA	Disk OK
	9	0.98 TB SATA	Disk OK
	10	1.95 TB SATA	Disk Unused
	11	(No Disk)	Empty Slot (No Disk)
	12	(No Disk)	Empty Slot (No Disk)
Node1234511 (X2, 12.7 TB)	1	0.98 TB SATA	Disk OK
	2	0.98 TB SATA	Disk OK
	3	1.95 TB SATA	Disk OK
	4	1.95 TB SATA	Disk OK
	5	0.98 TB SATA	Disk OK
	6	0.98 TB SATA	Disk OK
	7	1.95 TB SATA	Disk OK
	8	1.95 TB SATA	Disk OK
	9	0.98 TB SATA	Disk OK
	10	(No Disk)	Empty Slot (No Disk)
	11	(No Disk)	Empty Slot (No Disk)
	12	(No Disk)	Empty Slot (No Disk)

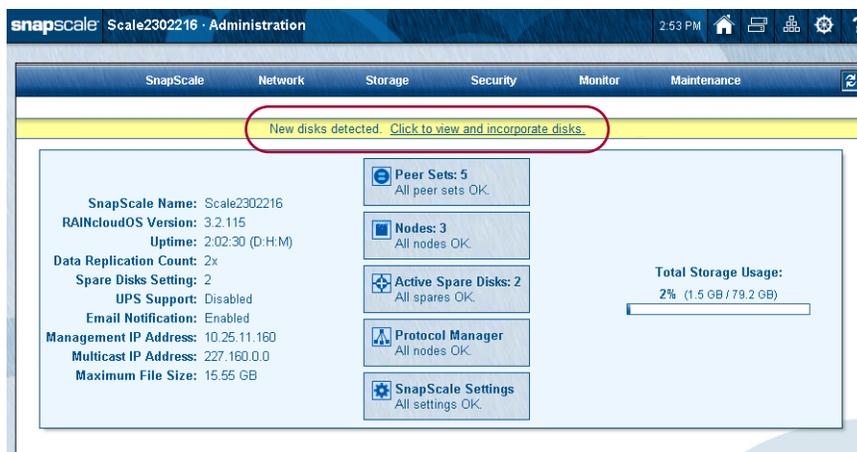
Legend: Disk OK Disk Unused Disk Failure Empty Slot (No Disk) Spare Disk Spare Too Small for Some Peer Sets Spare Too Small for Any Peer Set

[Click here to stop flashing LEDs on all nodes.](#)

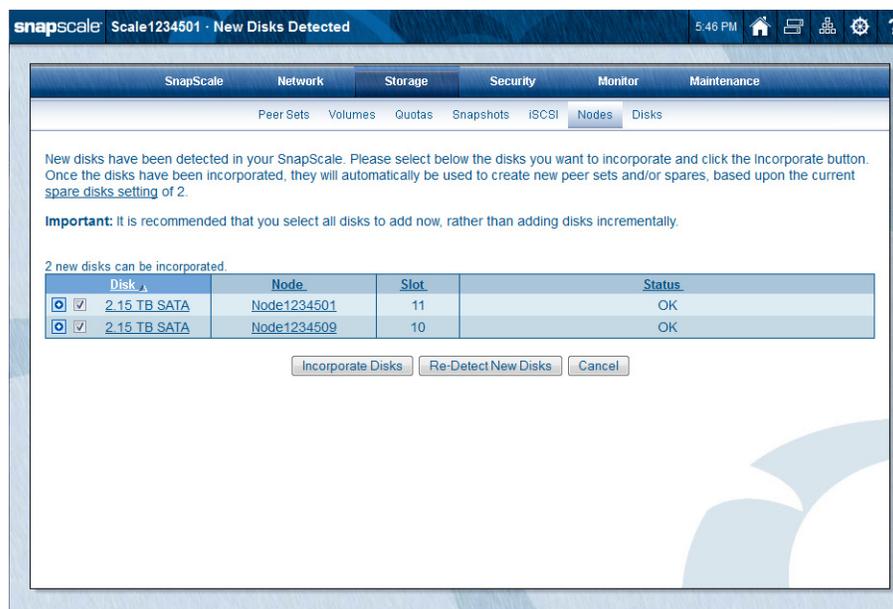
Refresh Close

Peer Set/Hot Spare Incorporation

When newly installed drives are detected, SnapScale first auto-incorporates drives to fix any failed peer set and fulfill any reserved spare count. The Web Management Interface then displays an alert banner about the new drives and the **New Disks Detected** page is activated.



Click the link in the alert to go to the **New Disks Detected** page. The page displays all new drives available to be formed into new peer sets or used as new spares.



The boxes next to the drive name can be unchecked to remove a drive from the incorporation list. However, to balance peer set and spare creation, it is recommended that all drives be incorporated at the same time.

Click **Incorporate Disks** to begin the process.

- If enough new or spare drives exist on different nodes based on your Data Replication Count (2x or 3x), new peer sets are formed as long as the spare count is satisfied.
- If there are not enough drives or they are not on different nodes, the drives are used to create additional hot spares.

After incorporation, the drives are displayed normally as peer sets or hot spares on the **Disks** page.

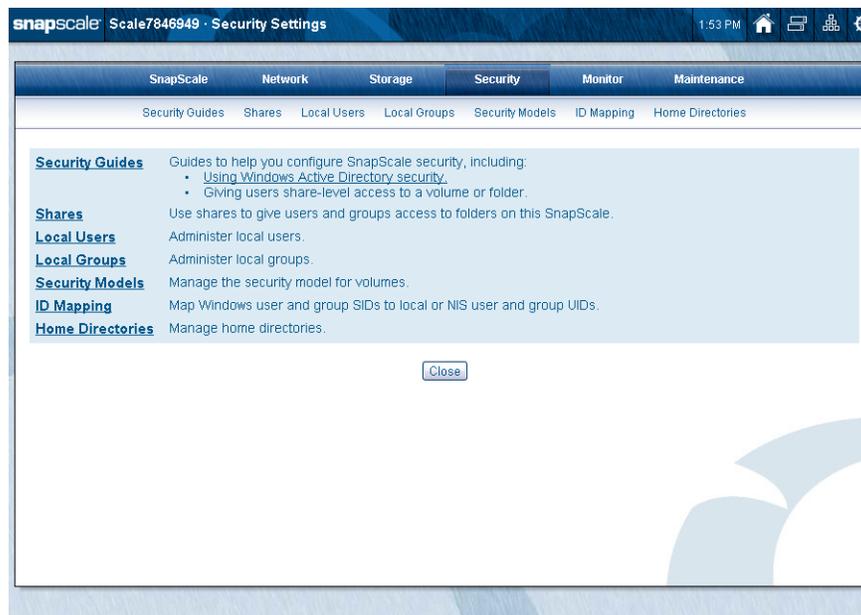
This section covers Security options for users, groups, shares, and file access.

Topics in Security Options

- [Overview](#)
- [Security Guides](#)
- [Shares](#)
- [Local Users](#)
- [Local Groups](#)
- [Security Models](#)
- [ID Mapping](#)
- [Home Directories](#)

Overview

Authentication validates a user’s identity by requiring the user to provide a registered login name (User ID) and corresponding password. SnapScale clusters have predefined local users and groups that allow administrative (admin) and guest user access to the cluster via all protocols. Those options are found on the **Security** tab:



Administrators may choose to join the SnapScale cluster to a Windows Active Directory domain, and CIFS/SMB clients can then authenticate to the cluster using their domain credentials. To accommodate NFS clients, the SnapScale cluster can also join an NIS domain, and can look up user IDs (UIDs) and group IDs (GIDs) maintained by the domain. See [User and Group ID Assignments](#) on page 6-3.

The SnapScale default security configuration provides one share to a default volume that can consume the entire cluster storage space. All network protocols for the share are enabled, and all users are granted read-write permission to the share via the guest account. By default, the guest user is disabled in SMB but enabled for HTTP.

Network clients can initially access the cluster using the guest account (where enabled), but if you require a higher degree of control over individual access to the filesystem for these clients, you must create local accounts (or use Windows Active Directory security for CIFS/SMB clients).

Local users or groups are created using **Security > Local Users** or **Security > Local Groups** in the Web Management Interface. Local users are also used for administrative access to the cluster through the cluster's Web Management Interface or SSM.

A local user or group is one that is defined locally on a SnapScale cluster using the Web Management Interface. The default users and groups listed below cannot be modified or deleted.

- **admin** – The local user admin account is used to log into the Web Management Interface. The default password for the admin account is also *admin*.
- **guest** – The local user guest account requires no password.
- **admingrp** – The Admin group account includes the default admin user account. Any local user accounts created with admin rights are also automatically added to this group.

Guidelines for Local Authentication

These password authentication guidelines are for both users and groups.

Duplicating Client Login Credentials for Local Users and Groups. To simplify user access for Windows Workgroup, duplicate their local client logon credentials on the SnapScale cluster by creating local accounts on the cluster that match those used to log on to client workstations. This strategy allows users to bypass the login procedure when accessing the cluster.



CAUTION: This strategy applies only to local users. Do not use duplicate domain user credentials if joined to an Active Directory domain.

Default Local Users and Groups. Default users and groups *admin*, *guest*, and *admingrp* appear on the list of users or groups on the user or group management pages, but they cannot be deleted or modified (although the admin password can be changed).

Changing Local UIDs or GIDs. The SnapScale cluster automatically assigns and manages UIDs and GIDs. Because you may need to assign a specific ID to a local user or group in order to match your existing UID/GID assignments, the cluster makes these fields editable.

Password Policies. To provide additional authentication security, set password character requirements, password expiration dates, and lockout rules for local users.

Local users can also be individually exempted from password expiration and character requirement policies. The built-in *admin* user is exempt from all password policies.

Local Account Management Tools. The following tools are available for creating, modifying, and editing local user and group accounts:

Function	Navigation Path
Local User Management	Navigate to the Local Users page, from which you can create, view, edit, and delete local users. You can also set user password policy, including password character requirements, maximum number of allowed logon failures, and password expiration settings.
Local Group Management	Navigate to the Local Groups page, from which you can create, view, edit, and delete local groups.

User and Group ID Assignments

A SnapScale cluster uses the POSIX standard to assign UIDs or GIDs, in which each user and group must have a unique ID. This requirement applies to all users and groups on the cluster, including NIS, Windows domain, and local users plus NIS groups.

If you join the cluster to a Windows or NIS domain, IDs are assigned using available IDs only. Consider the following when creating users and groups:

- UIDs and GIDs from 0 to 100 are unavailable for use. If you try to assign a UID or GID that is less than 101 (or in use by NIS or the Windows domain), you will get an error message.
- When the cluster automatically generates UIDs or GIDs for imported Windows domain users or groups, UIDs or GIDs that are already in use by local and NIS users are skipped.
- When NIS domain users and groups are imported, the cluster discards any UIDs that are less than 101 or are in conflict with UIDs already in use by local or Windows domain users and groups.

The `nfsnobody` and `nobody` user IDs (UID 65534 and 65535, respectively) and GIDs are reserved. They are not mappable to other IDs, nor is another ID mappable to `nfsnobody` or `nobody`.

Security Guides

Security Guides are special wizards to guide you through:

- Setting up Windows Active Directory security.
- Giving users or groups share-level access to an entire volume.
- Giving users or groups share-level access to a folder on a volume.



Windows Active Directory Security Guide

The **Windows Active Directory Security Guide** wizard guides you through the setup of Windows Active Directory on your cluster.

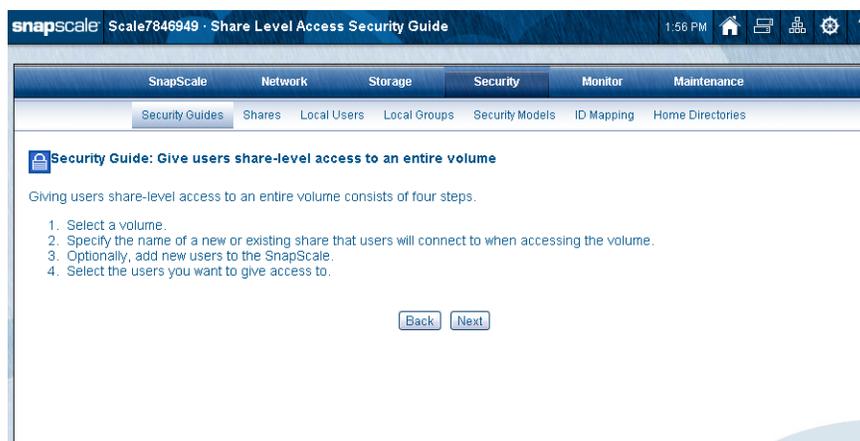
NOTE: You cannot join an Active Directory domain if NTP is enabled. If you see such a message, click the NTP link to change your settings.

When the cluster joins a domain, it does so as a single unit under the cluster name, and all nodes operate equally under the cluster name to authenticate against the domain. This provides multipoint domain-authenticated access to the cluster through each node.



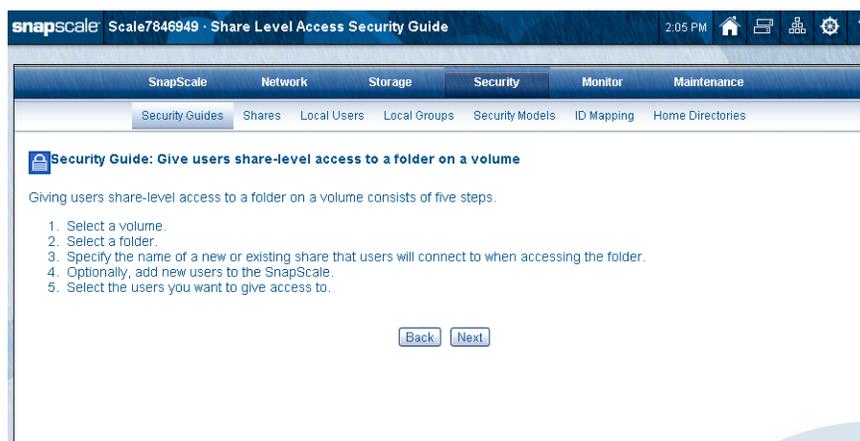
Entire Volume Security Guide

This **Share Level Access Security Guide** wizard guides you through the four steps it takes to give share-level access to an entire volume.



Folder on Volume Security Guide

This **Share Level Access Security Guide** wizard guides you through the five steps it takes to give share-level access to a folder on a volume.



Shares

SnapScale provides full integration with existing Windows Active Directory domain or UNIX NIS user and group databases. At the share level, administrators can assign read-write or read-only share access to individual Windows (and local) users and groups. Administrators can also edit the NFS exports file to control how shares are exported to NFS client machines.

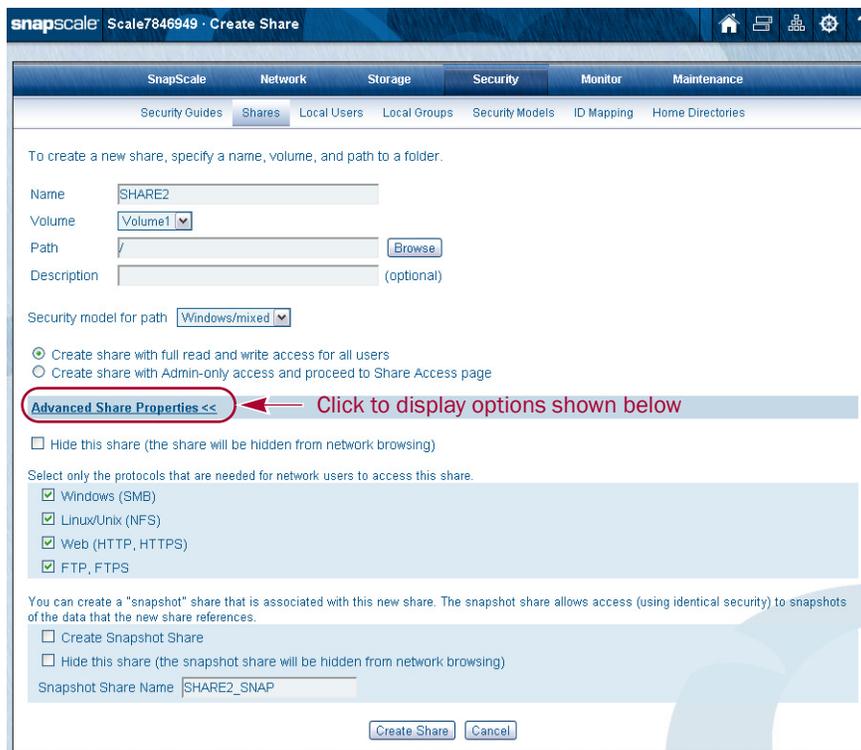


Share Security Overview

SnapScale supports file access in Windows and UNIX networks. New shares are created by default with full read-write access to all users, subject to the filesystem permissions on the share target directory. The first step to securing a cluster is to specify access at the individual share level. Administrators can assign read-write or read-only share access to individual Windows (and local) users and groups.

Create Shares

To create a new share, you need, at a minimum, to specify the share name, volume, and folder path. Click **Create Share** on the default **Shares** page to start the process.



NOTE: The snapshot information at the bottom is only shown if snapshot space has been reserved.

By clicking the **Advanced Share Properties** link, additional options are displayed. Use these options to hide the share from network browsing, select the protocols supported, and create a snapshot share associated with this share.

Creating a Share

Creating a share includes selecting the volume, security model, and directory path for the share and then defining share attributes and network access protocols.

1. Accept the default **share name** or enter a new one.
To ensure compatibility with all protocols, share names are limited to 27 alphanumeric characters (including spaces).
2. Choose the **volume** you need from the drop-down menu.
3. Select from the following **path options**:
 - **To create a share to the entire volume** – The current Path field defaults to the root path of the volume. Simply leave it blank if this is the desired configuration.
 - **To create a share to a folder on the volume** – Browse to the folder to which you want to point the share, click the folder name, and click **OK**.

NOTE: If you want to create a new folder inside any other folder, type the folder name into New Folder Name and click Create Folder.

4. If desired, enter a **description** to clarify the purpose of the share.
5. Choose a **security model for path** by selecting either **Windows/Mixed** or **UNIX** from the drop-down list.
The option defaults to the current security model at the specified path. If changed to a different security model, the change will propagate to all files and subdirectories underneath. For more information, see [Security Models](#) on page 6-25.
6. Choose the user-based **share access** option desired.
Choose either **Create share with full read and write access for all users**, or **Create share with Admin-only access and proceed to Share Access page** to configure the user share access. For more information, see [Share Access Behaviors](#) on page 6-10.

NOTE: If selecting the share with Admin-only access option and the share has NFS enabled, be sure to configure the NFS access settings afterward.

7. To further configure the share, click **Advanced Share Properties**, and enter any of the following:

Option	Description
Hide this Share	Select this option if you want the share to be hidden from network browsing using SMB and HTTP/HTTPS protocols (but not NFS).
Protocols	Select the access protocols for the share: Windows (SMB), Linux/UNIX (NFS), or Web (HTTP/HTTPS).
Snapshot Share	To create a snapshot share, check the Create Snapshot Share box. Optionally, do either of the following: <ul style="list-style-type: none"> • To hide the snapshot share from the SMB and HTTP protocols, check the Hide Snapshot Share box. • If you do not want to accept the default name provided, enter a unique name for the Snapshot Share Name field. Use up to 27 alphanumeric characters (including hyphens and spaces).

8. Click **Create Share** to complete the process.

Edit Share Properties

NOTE: You cannot change the volume (or path) of a share once it is created. If you need to change the volume, you must delete the share and create a new one on the other volume.

Once a share has been created, you can change its name, description and the advanced properties. To edit the properties, go to **Security > Shares > Share Properties** (displayed by clicking the share name in the table).

By clicking the **Advanced Share Properties** link, additional options are displayed. Use these options to hide the share from network browsing, select the protocols supported, and create a snapshot share associated with this share.

NOTE: The snapshot information at the bottom is only shown if snapshot space has been reserved.

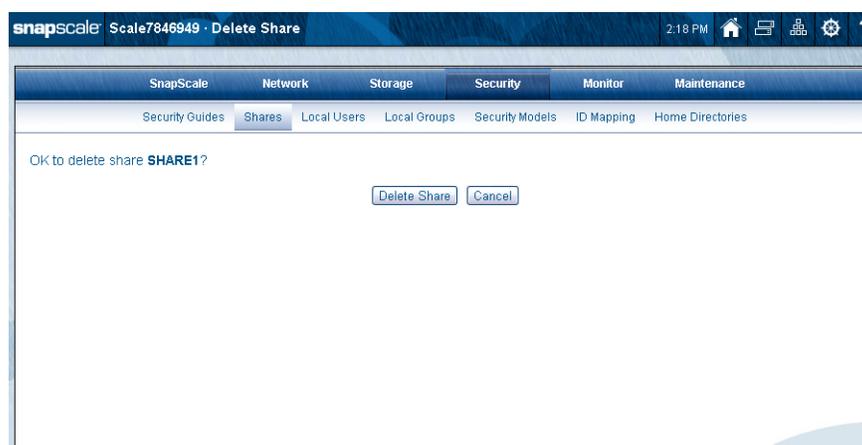
Option	Description
Name	Accept the default share name or enter a new one. If you change the default, observe the following guidelines: <ul style="list-style-type: none"> • Make sure the share name is unique on this cluster. • To ensure compatibility with all protocols, share names are limited to 27 alphanumeric characters (including hyphens and spaces).
Description	If desired, enter a description of the share. This is an opportunity to clarify the purpose of the share.

Option	Description
Hide this share	Select this option if you want the share to be hidden from network browsing using SMB and HTTP/HTTPS (but not NFS) protocols.
Protocols	Select the access protocols for the share: Windows (SMB), Linux/UNIX (NFS), or Web (HTTP/HTTPS).
Snapshot Share	<p>The option that displays depends on whether a snapshot share currently exists.</p> <p>To create a snapshot share, check the Create Snapshot Share box. Optionally, do either of the following:</p> <ul style="list-style-type: none"> To hide the snapshot share from the SMB and HTTP protocols (but not NFS), check the Hide Snapshot Share box. If you do not want to accept the default name provided, enter a unique name for the Snapshot Share Name field. Use up to 27 alphanumeric characters (including hyphens and spaces). <p>To remove a existing snapshot share, check the Remove Snapshot Share box.</p>

Delete Shares

To delete a share, go to **Security > Shares > Share Properties** (displayed by clicking the share name in the table).

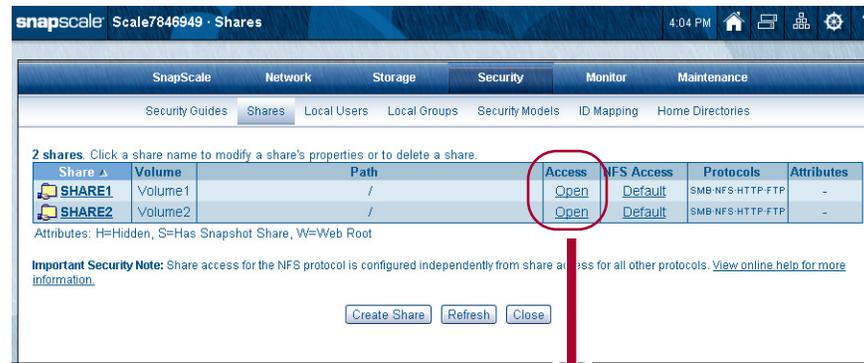
1. At the **Delete Share** page, click **Delete Share**.
2. At the confirmation page, click the **Delete Share** button again.



Configuring Share Access

In **Security > Shares**, in the **Access** column, click the link next to the share you want to configure. The **Share Access** page is displayed. You can set access levels for the share, as well as grant or deny access to specific users and groups.

NOTE: To add a new user to a share, you must first create the user, then add that user to the share. Please see Local Users on page 6-16 for information on creating new users.



Share Access Behaviors

Administrators tasked with devising security policies for SnapScale clusters will find the following share access behaviors informative:

- **Share access defaults to full control** – The default permission granted to users and groups when they are granted access to the share is full control. You may restrict selected users and groups to read-only access.
- **User-based share access permissions are cumulative** – An SMB or HTTP user's effective permissions for a resource are the sum of the permissions that you assign to the individual user account and to all of the groups to which the user belongs in the **Share Access** page. For example, if a user has read-only permission to the share, but is also a member of a group that has been given full-access permission to the share, the user gets full access to the share.
- **NFS access permissions are not cumulative** – An NFS user's access level is based on the permission in the NFS access list that most specifically applies. For example, if a user connects to a share over NFS from IP address 192.168.0.1, and the NFS access for the share gives read-write access to "*" (All NFS clients) and read-only access to 192.168.0.1, the user will get read-only access.

- **Interaction between share-level and file-level access permissions** – When both share-level and file-level permissions apply to a user action, the more restrictive of the two applies. Consider the following examples:

Example A: More restrictive file-level access is given precedence over more permissive share-level access.

Share Level	File Level	Result
Full control	Read-only to File A	Full control over all directories and files in SHARE1 <i>except</i> where a more restrictive file-level permission applies. The user has read-only access to File A.

Example B: More restrictive share-level access is given precedence over more permissive file-level access.

Share Level	File Level	Result
Read-only	Full control to File B	Read-only access to all directories and files in SHARE1, <i>including</i> where a less restrictive file-level permission applies. The user has read-only access to File B.

Setting User-based Share Access Permissions

Share permissions for Windows and HTTP users are configured from **Security > Shares** by clicking the link in the **Access** column of the share you want to configure. Share permissions for NFS are configured and enforced independently. See [NFS Access for Shares](#) on page 6-13 for more information.



User-based share access permissions apply to users connecting over SMB or HTTP. Users and groups with assigned share access permissions appear in the list on the left (**Users and groups with specific access to share**). To search for those without assigned access, use the box on the right (**Search for users and groups**).

The default permission granted to users and groups when they are granted access to the share is **Full Access**. You may restrict selected users and groups to **Read-only Access**.

Share-Level Access Permissions	
Full	Users can read, write, modify, create, or delete files and folders within the share.
Read-only	Users can navigate the share directory structure and view files.

1. Display the **Share Access** page (**Security > Shares > *access_link***).
2. To **add** share access permissions for a user or group:
 - a. At the bottom, using the drop-down list, select the **domain or local user/group list** to search.

NOTE: For domains that require authentication (showing an "(A)" after the name), after selecting the domain name, enter the User Name and Password for that domain. The user name and password can be for any user in the domain and are used to retrieve basic information (like the user & group lists) from the domain.

- b. Enter the **search string** (or select **Find All**).

When entering a search string:

- Returned results will include all users and groups whose name **begins** with the string entered in the Search field.
- The search results returned may be limited. Fine tune your search by using a more specific string to return the names desired.
- On the rare occasion you need to search for a domain that is not listed ("remote domain"), select a domain from the Search drop-down list through which to search, then enter in the Find box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, `remote_domain\user_name`).

- c. Click **Search** to display any matches.

After you click **Search**, another authentication prompt may be presented to authenticate with the remote domain.

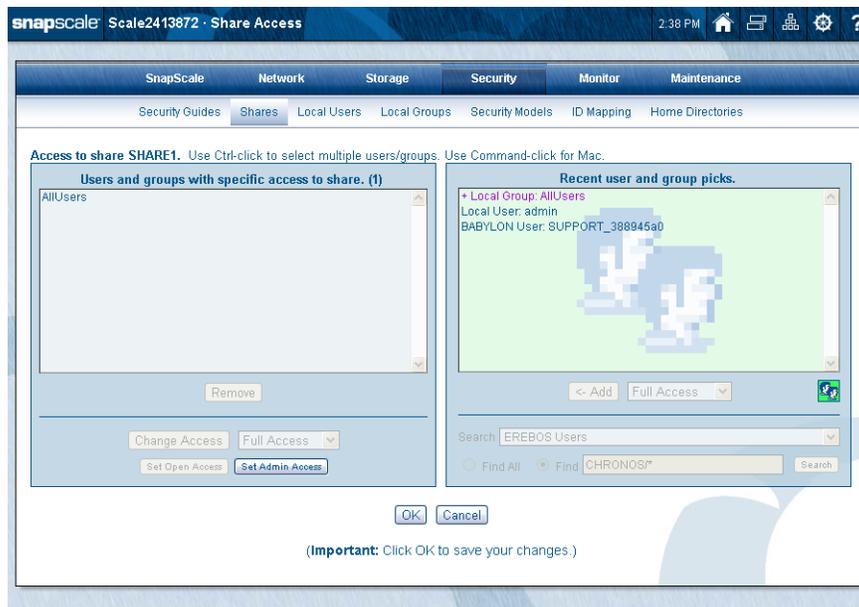
- d. Select one or more **names** in the list.

Users that already have access are shown in purple font with a plus sign (+) in front of their name.

- e. Choose either **Full Access** or **Read Only** from the drop-down list.

- f. Click **Add**.

NOTE: To display recent user or group picks, click the faces (👤) icon. A list with a green background is displayed. Click the now green icon to return to the normal search box.



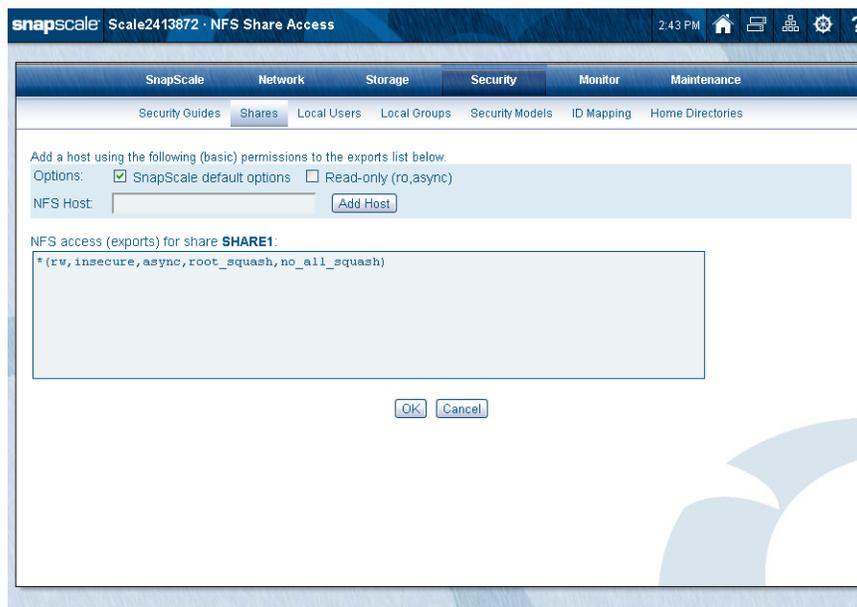
3. To **remove** share access permissions for a user or group:
 - a. Select one or more **users or groups** in the left box.
 - b. Click **Remove**.
4. To change **access permissions** for a user or group, select one or more users or groups in the left box, then select either **Full Access** or **Read Only** from the drop-down list, and click the **Change Access** button.
5. To quickly specify either Open or Admin-only **access** for the entire share, click either the **Set Open Access** or **Set Admin Access** button.
6. Click **OK** to save share permissions.

NFS Access for Shares

NOTE: Multiple shares pointing to the same target directory must have the same NFS access settings. The Web Management Interface applies the same NFS access for all shares pointing to the same directory.

To configure NFS access, click the link shown in the **NFS Access** column for the share you want to configure. You can configure NFS access to the share using standard Linux “exports” file syntax.

On the **Shares** page, click the name of the access type listed in the **NFS Access** column to open the **NFS Share Access** page.



The NFS access text box is a window into the client access entries in the cluster's *exports* file. This file serves as the access control list for filesystems that may be exported to NFS clients. You can use the **Add Host** controls as described below to assist in making entries to the file, or you can directly edit the text box. After all entries are made, click **OK** to return to the **Shares** page.

NOTE: The syntax used in this file is equivalent to standard Linux exports file syntax. If the cluster detects any errors in syntax, a warning message appears. You can choose to correct or ignore the error warning.

The Exports File Default Options. The SnapScale default setting provides read-write access to all NFS clients.

```
*(rw,insecure,async,root_squash,no_all_squash)
```

The entry options are explained in the following table:

Entry Code	Meaning
Asterisk	All NFS clients
ro	The directory is shared read only (ro).
rw	The client machine will have read and write (rw) access to the directory.
insecure	Turns off the options that require requests to originate on an Internet port less than IPPORT_RESERVED (1024).
root_squash	Forces users connected as root to interact as the "nobody" user (UID 65534). This is the RAINcloudOS default.

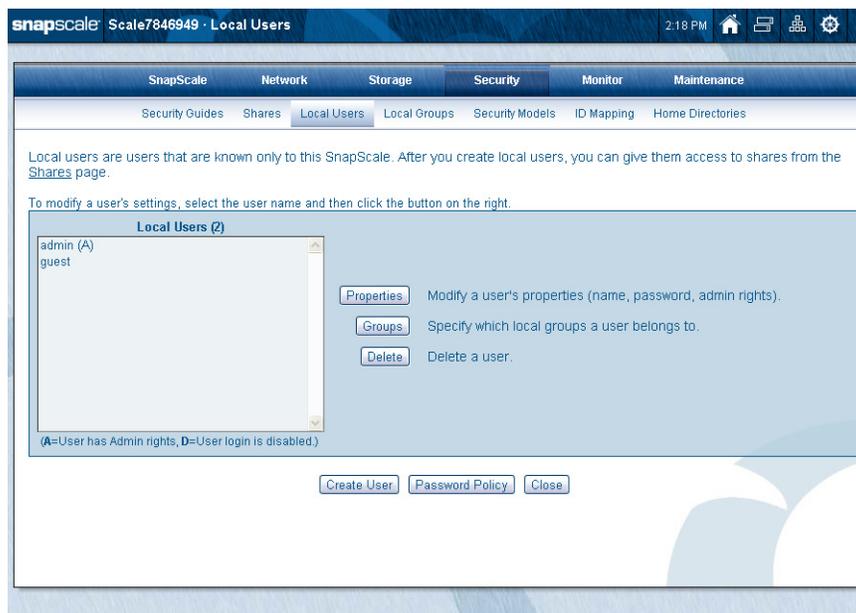
Entry Code	Meaning
no_root_squash	no_root_squash means that if root is logged in on your client machine, it will have root privileges over the exported filesystem. By default, any file request made by user root on the client machine is treated as if it is made by user nobody on the cluster. (Exactly which UID the request is mapped to depends on the UID of user nobody on the cluster, not the client.) If no_root_squash is selected, then root on the client machine will have the same level of access to the files on the system as root on the cluster. This can have serious security implications, although it may be necessary if you want to perform any administrative work on the client machine that involves the exported directories. You should not specify this option without a good reason.
async	Tells a client machine that a file write is complete – that is, has been written to stable storage – when NFS has finished handing the write over to the filesystem.
no_all_squash	Allows non-root users to access the nfs export with their own privileges.

Using the Add Host Option. Follow these steps:

1. Select **one** of the following options:
 - **SnapScale Default Options** – Inserts the default options as described above
 - **Read Only** – Inserts the read only option only
 - **Both** – Inserts default options, but substitutes read only for read/write
2. Do **one** of the following in the NFS host text box:
 - **To apply the options to all NFS hosts** – Leave this field blank
 - **To apply the options to specific hosts** – Enter one or more IP addresses.
3. Click **Add Host**.

Local Users

The **Local Users** page provides all the options to manage local users. Local users are users that are known only to the cluster being accessed. Each SnapScale cluster comes with two predefined users: admin and guest. The admin user has full Administrator rights. Go to **Security > Local Users** to view settings or make changes.



Create a User

Click the **Create** button to create a new user on this cluster. Enter the user data, select any special options, and click the **Create User** button again.



To Create a Local User

1. On the **Local Users** page, click **Create User**.
2. On the **Create Local User** page that opens, enter the requested **information**:

Option	Description
Name	Use up to 31 alphanumeric characters and the underscore.
Full Name	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
Password	Passwords are case-sensitive. Use up to 15 alphanumeric characters without spaces.
Confirm Password	Type the chosen password again for verification.
User ID (UID)	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see User and Group ID Assignments on page 6-3 .
Disable User Login	Check this box to disable the user login. The user's information will remain in the system, but login rights are denied. The user login can be re-enabled by clearing the box. This box can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.
Exempt from Password Expiration and Character Requirements	This box is only visible if Password Policy is enabled. Check this box to exempt this user from password expiration and character requirement policies.
Grant Admin Rights To This User	Check this box to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click **Create User** again to create the user account.

Edit User Properties

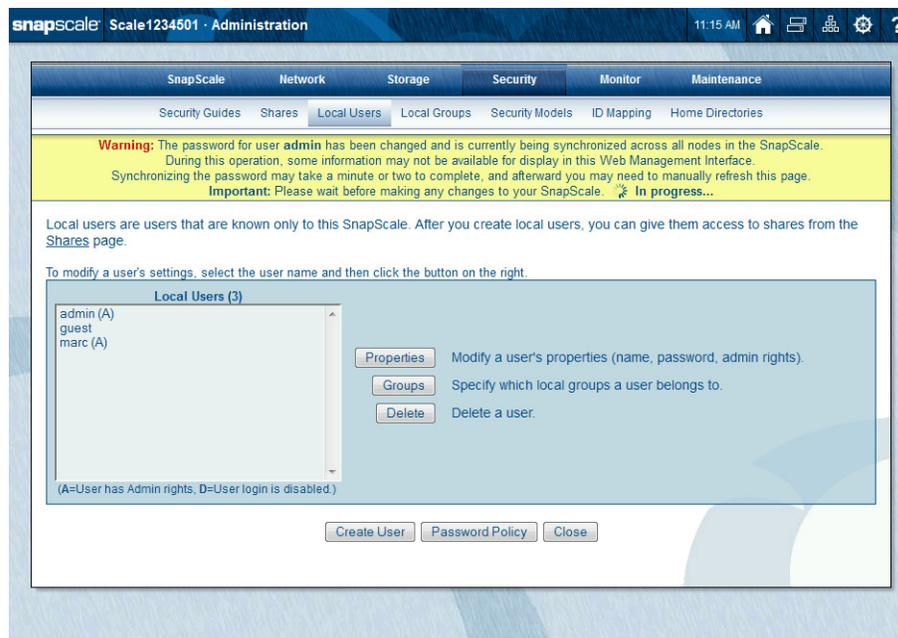
Use the **Properties** button to open the **Local User Properties** page to make changes to the user's full name, password, or user ID (UID). Note that the UID cannot be changed for the built-in admin user.

The screenshot shows the 'Local User Properties' dialog box in the SnapScale web management interface. The dialog has a title bar with 'snap scale Scale7846949 · Local User Properties' and a system tray showing '2:20 PM'. The main content area has a navigation bar with 'SnapScale', 'Network', 'Storage', 'Security', 'Monitor', and 'Maintenance'. Below this is a sub-navigation bar with 'Security Guides', 'Shares', 'Local Users', 'Local Groups', 'Security Models', 'ID Mapping', and 'Home Directories'. The 'Local Users' section is active, showing the following fields:

- Name: admin
- Full Name: admin (optional)
- Password: (empty field) (Leave blank to keep existing password)
- Confirm Password: (empty field)
- User ID (UID): 1

 At the bottom, there is a checkbox labeled 'Grant admin rights to this user (A local user with admin rights will be able to access this Web Management Interface.)' which is checked. Below the checkbox are 'OK' and 'Cancel' buttons.

IMPORTANT: When changing the Admin password, it can take a minute or so to synchronize the new password across all nodes. During this time, a warning message is displayed in the Web Management Interface. While the cluster, all the nodes, and all the data are fully accessible during this synchronization process, you should wait for the message to disappear before making further changes to your SnapScale.



To Edit Local User Properties

1. On the **Local Users** page, highlight the user you want to edit and click **Properties**.
2. On the **Local User Properties** page that opens, enter or change any of the **information**:

Option	Description
Name	NOTE: Cannot be modified.
Full Name	Use up to 49 alphanumeric characters (includes spaces). Input in this field is optional.
Password	Passwords are case-sensitive. Leave this field blank to keep the existing password.
Password Verify	Type the chosen password again for verification. Leave this field blank to keep the existing password.
User ID (UID)	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see User and Group ID Assignments on page 6-3. NOTE: Changing a user's UID may alter filesystem access permissions that apply to that UID. In addition, any existing permissions for a UID previously assigned to a user that are changed to a different UID may become active if another user is created with the same UID. Carefully consider security configuration on existing files and directories before changing the UID of a user.

Option	Description
Disable User Login	Check this box to disable the user login. The user's information will remain in the system, but login rights are denied. The user login can be re-enabled by clearing the box. This box can also be used to enable a user locked out by the <i>Disable login after n attempts</i> password policy.
Exempt from Password Expiration and Character Requirements	NOTE: This box is only visible if Password Policy is enabled. Check this box to exempt this user from password expiration and character requirement policies.
Grant Admin Rights To This User	Check this box to allow the user access to the Web Management Interface and SSH (for access to the CLI and backup agent installation).

3. Click **OK**.

Local User Password Policies

NOTE: Local users can be individually exempted from password expiration and character requirements. This may be necessary for some special users, such as users configured to perform backups. See [To Create a Local User](#) on page 6-16 for procedures to set password policy for local users. Also, the built-in *admin* user is automatically exempt from all password policies.

Use the **Password Policy** button to make changes to the local user password settings.

To Set Password Policy for Local Users

1. On the **Local Users** page, click the **Password Policy** button.
2. On the **Local Users Password Policy** page, check the **Enable Password Policy** box.

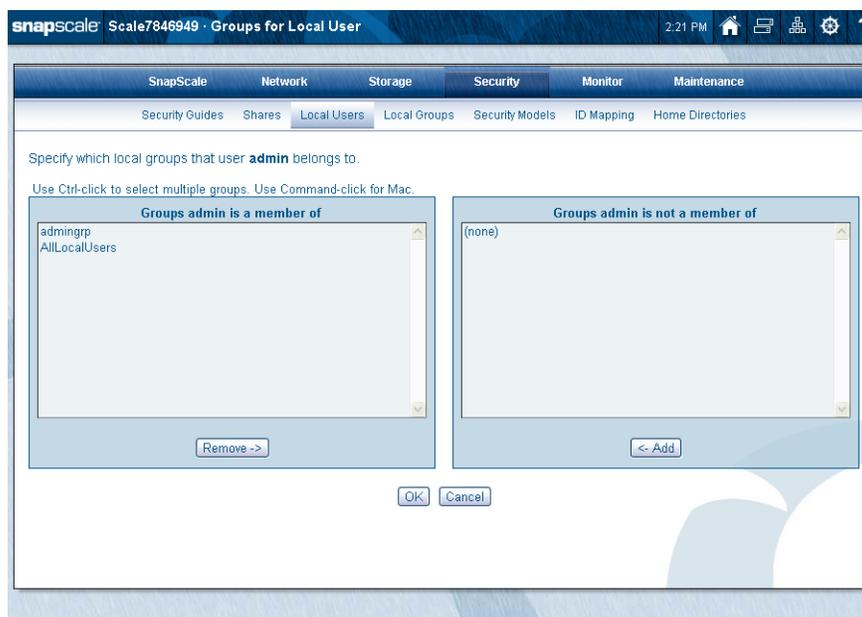
3. Enter the following **information**:

Option	Description
Character Requirements	Select the alpha/numeric/special character requirements for the password from the drop-down list.
Minimum Number of Characters	Check this box to enable the policy, then enter the minimum number of characters required for the password.
Disable Login After <i>n</i> Attempts	Check this box to enable the policy, then enter the number of times a user can fail to login before the system locks the user out. This applies to failed logins when connecting to any node in the cluster. NOTE: To unlock a user, clear the Disable User Login box for the user in the Local Users page.
Re-enable a Disabled Login After <i>n</i> Minutes	If you have defined a limit to the number of times a user can fail to log in, you can also check this box and enter a time period after which the system will allow the user to log in again. NOTE: This saves the administrator from having to manually re-enable the user.
Expire Password After <i>n</i> Days	Check this box to enable the policy, then enter the number of days before the password must be changed. NOTE: Local users with expired passwords can change their passwords at: <a href="http://<clustername>/changepassword">http://<clustername>/changepassword .

4. Click **OK** to save the settings.

Assign User to Group

Use the **Groups for Local User** page (**Security > Local Users > Groups**) to make changes to a local group membership.



To Add or Remove Users from Groups

1. On the **Local User** page, select a **user**.

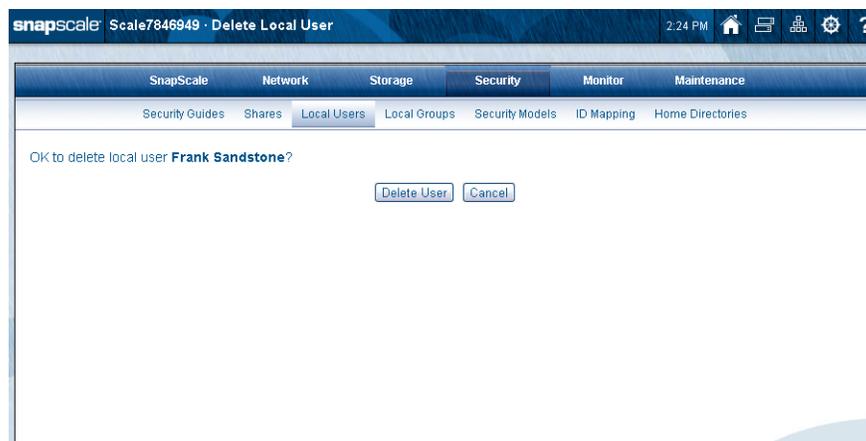
2. Click **Groups**.
The group settings for the selected user are shown.
3. To add the selected user to a group, select the group from the right-side list and click **<- Add**.
4. To delete the selected user from a group, select the group from the left-side list and click **Remove ->**.
5. Click **OK** to save your changes.

Delete Local User

On the **Local Users** page, use the following process to remove a user.

To Delete a Local User

1. On the **Local Users** page, select the user to be deleted.
2. Click **Delete**.
The confirmation page is displayed.
3. Click **Delete User** to delete the selected user (or click **Cancel**).



Local Groups

The **Local Groups** page (**Security > Local Groups**) provides all the options to manage local groups. Local groups are groups of local users that are known only to the cluster being accessed. Each SnapScale cluster comes with one predefined group: `admingrp`.



Create New Group

Use the **Create** button to create a new group on this cluster. Options include the group name and changing the Group ID (GID).



To Create a New Local Group

1. On the **Local Groups** page, click **Create Group**.
2. On the **Create Local Group** page that opens, enter the following **information**:

Option	Description
Group Name	Use up to 31 alphanumeric characters and the underscore.

Option	Description
Group ID (GID)	Displays the user identification number assigned to this user. Alter as necessary. For information on available UID ranges, see User and Group ID Assignments on page 6-3 .

3. Click **Create Group** when finished.
4. The **Users for Local Group** page is displayed, allowing you to immediately add users to your new group.
5. Click **OK** when you are finished adding users.

Edit Group Properties

Use the **Properties** button to open the **Local Group Properties** page to make changes to the options there.



To Edit Local Group Properties

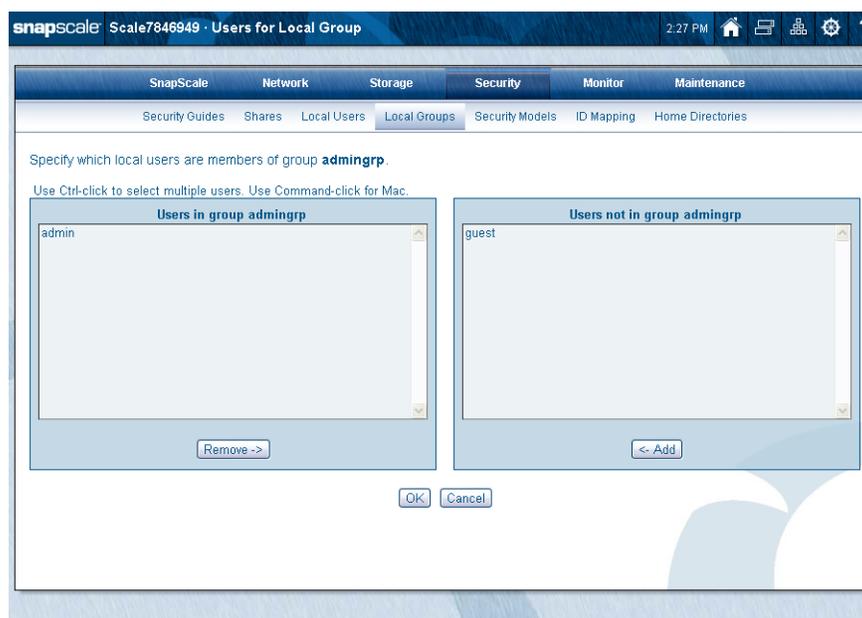
1. On the **Local Groups** page, select the group you want to edit and click **Properties**.
2. On the **Local Groups Properties** page that opens, you can change the **GID**. For information on available UID ranges, see [User and Group ID Assignments](#) on page 6-3.

NOTE: Changing a group's GID may alter filesystem access permissions that apply to that GID. In addition, any existing permissions for a GID previously assigned to a group that are changed to a different GID may become active if another group is created with the same GID. Carefully consider security configuration on existing files and directories before changing the GID of a group.

3. Click **OK**.

Specify Users in Group

Use the **Users for Local Group** page (**Security > Local Groups > Users**) to make changes to a local group membership.



To Add or Remove Group Users

1. On the **Local Groups** page, select a group name and click **Users**.
2. Add users by selecting the user in the right-side list and clicking **<- Add**.
3. Delete users by selecting the user in the left-side list and clicking **Remove ->**.
4. Click **OK** when finished.

Delete Group

1. On the **Local Groups** page, select the group to be deleted and click **Delete**. The confirmation page is displayed.
2. Click **Yes** to delete the selected group, or click **No** to cancel the deletion.

Security Models

There are two file-level security models that can be used by a SnapScale cluster: Windows/Mixed and UNIX. The security model can only be configured on the volumes.

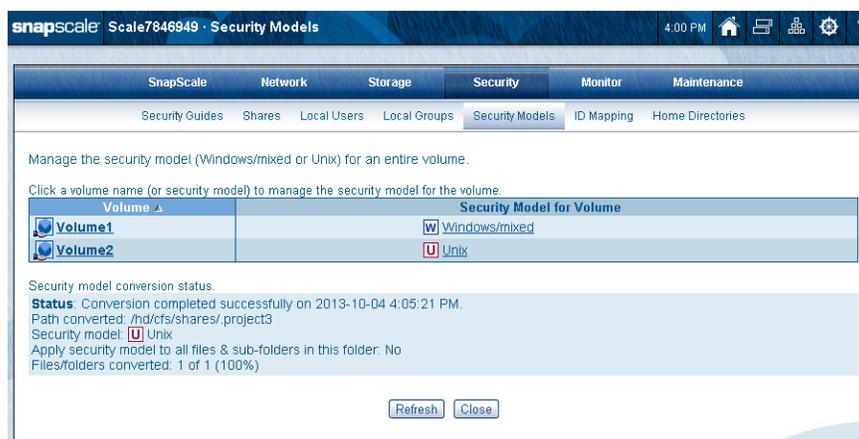
The security model determines the rules regarding which security personality that is present on files and folders created by the various protocols and clients, and whether the personality of files and folders can be changed by changing permissions.

Folders created in a volume default to the security model of that volume. The folder's security model may differ from the personality of the folders (for example, folders with a Windows/Mixed security may have a UNIX personality).

For more information about security models, see [Appendix A, Backup Solutions](#).

Managing Volume Security Models

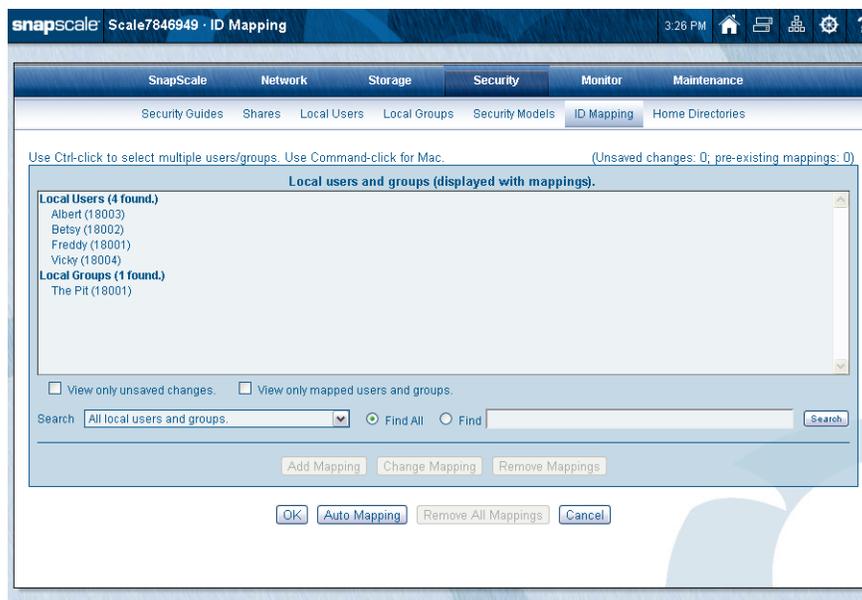
1. Select **Security > Security Models**.



2. Click the **security model name** (Windows/Mixed or UNIX).
Clicking the **Volume** name does the same thing.
3. From the drop-down list, select the **security model type** desired, and click **OK**.
4. At the confirmation message, click **Apply Security Model**.
If there are files and directories under the volume, you are prompted whether you want to recursively apply the change. When done, the main page displays a conversion status.

ID Mapping

ID mapping allows users and groups that exist on Windows domains to share user IDs with local or NIS users and groups. This results in the same permissions and quota consumption applying to both the Windows domain user and the local or NIS user.



Select a local or NIS user or group from the displayed list on the default page. You can then use **Add Mapping** to map the user's UID or group's GID to that of a Windows domain user or group. **Change Mapping** is used to change existing mappings. **Remove Mappings** removes one or more mappings while **Remove All Mappings** removes all mappings that had been previously established.

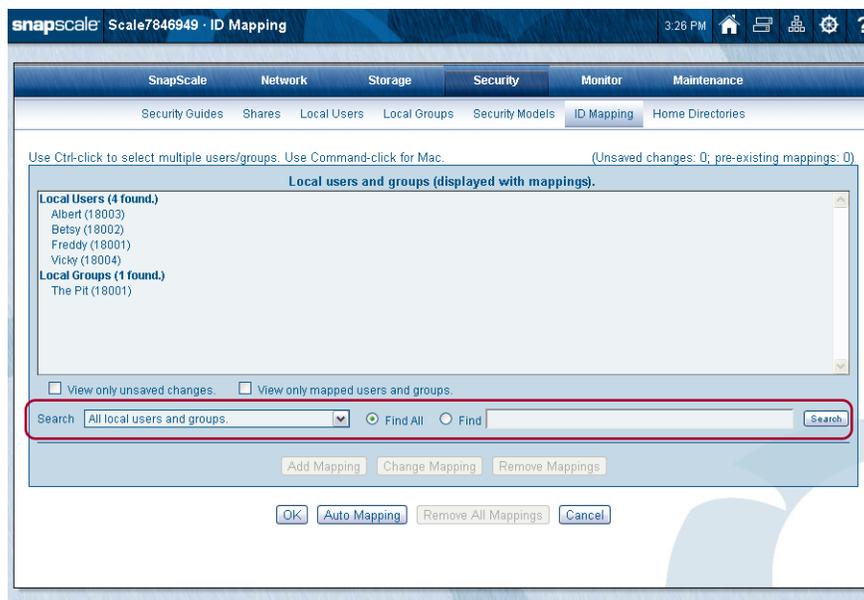
Options to simplify the discovery of a desired user or group to manage their ID mapping search options are presented at the bottom of the selection pages:

- Check **View only unsaved changes** to display only mapping changes that have not yet been applied.
- Check **View only mapped users and groups** to display only local or NIS users and groups that have been mapped to a Windows domain user or group.

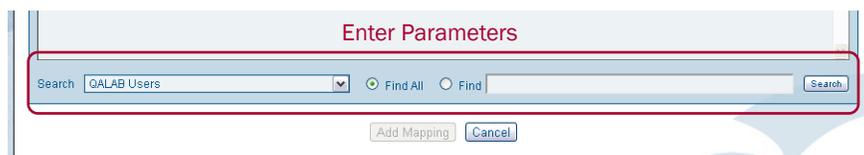
Add Mapping

Follow this procedure to map a user or group:

1. If the desired user or group to be mapped to does not appear in the **ID Mapping** page list, use the **search option** to locate it.



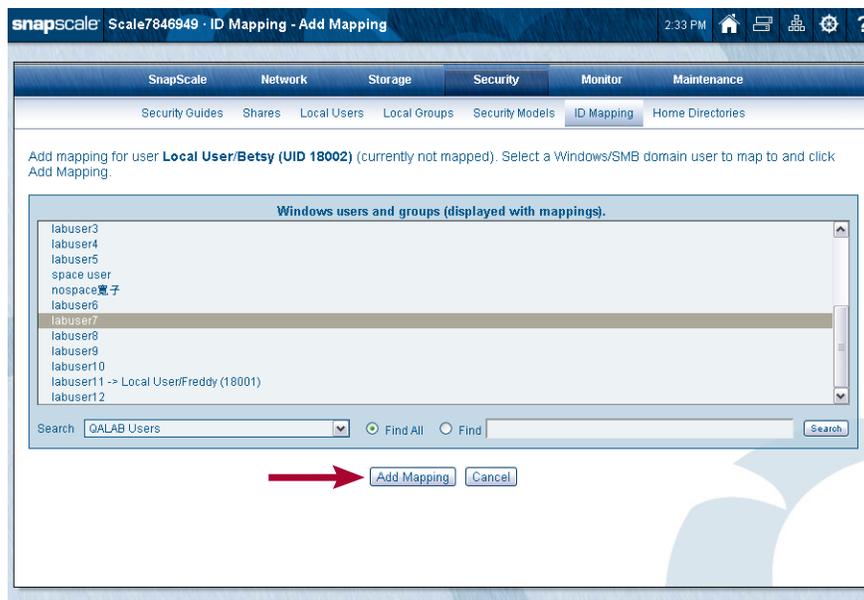
- a. At the bottom of the list, using the **Search** drop-down list, select the local or NIS user or group **list** to be searched.
 - b. Select **Find** and enter the **search string** (or select **Find All**).
Enter the exact **name** (or a string with a wildcard "*" before or after) as the search string.
 - c. Click **Search** to display any matches.
2. Select a **user or group** from the results list, and click **Add Mapping**.
 3. At the **Add Mapping** page, to find the user/group you want to map to, select the Windows domain **user or group list**, the search scope, enter a search string if needed, and click **Search**.
 - To search for a specific user or group, use either **Find All** or a **Find** search string (wildcard "*" before or after is allowed).



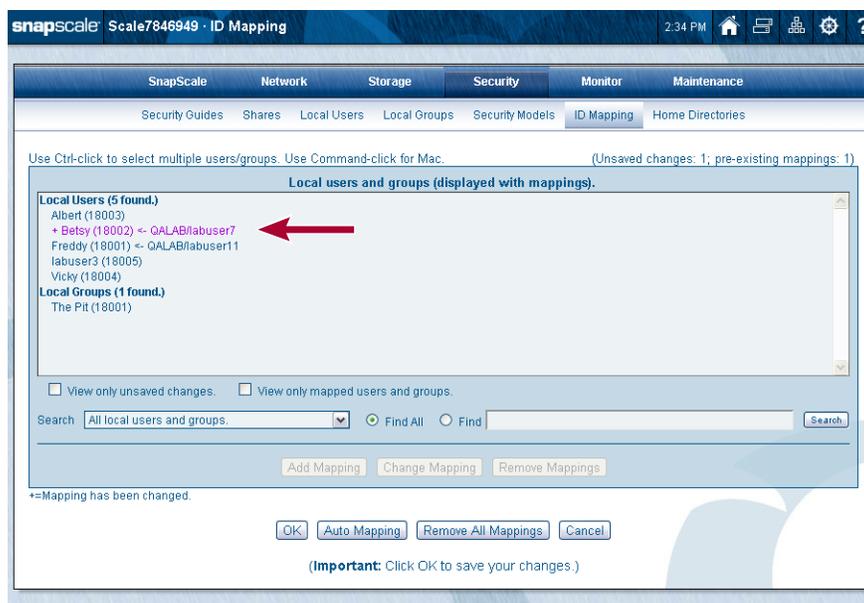
- For domains that **REQUIRE** authentication (showing an "(A)" after the name), select the domain name, enter the user name and password for that domain, and use either **Find All** or a **Find** search string (using the first few letters of the user/group name).



- On the rare occasion you need to search for a Windows domain that's not listed (“remote domain”), select a Windows domain from the **Search** drop-down list through which to search, then enter in the **Find** box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, `remote_domain\user_name`). After you click **Search**, another authentication prompt may be presented to authenticate with the remote domain.
4. From the search results, select the Windows domain **user/group** to which you want to map the local or NIS user, and click **Add Mapping**.



The mapping result is shown on the default page with the users/groups that were mapped in purple with a plus (+) in front of their name.



5. Repeat **Steps 1–4** to add **additional mappings**.
6. Click **OK** to save changes (or **Cancel** to reset).

7. When done with all your mappings, click **OK** to activate them.
8. At the confirmation page, click **Save Changes**.
9. At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

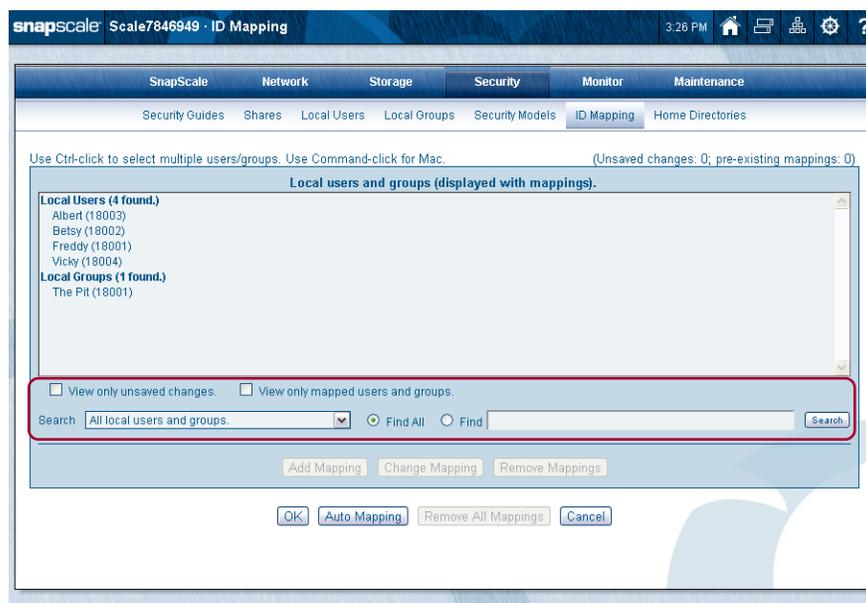
 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

See [Filesystem Updates](#) on page 6-37 for more details.

Change Mapping

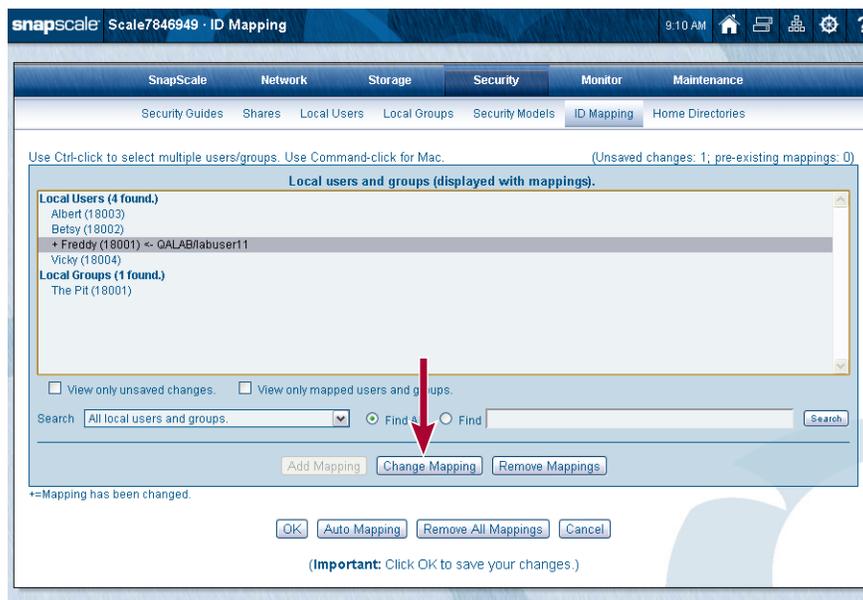
To map an already mapped local or NIS user or group to a different Windows domain user or group, follow these steps:

1. If the desired user or group to be changed does not appear in the default page list, use the **search option** to locate them.



- a. At the bottom of the list, using the **Search** drop-down list, select the local or NIS user or group **list** to be searched.
- b. Select **Find** and enter the **search string** (or select **Find All**).
Enter the exact **name** (or a string with a wildcard “*” before or after) as the search string.
- c. Click **Search** to display any matches.

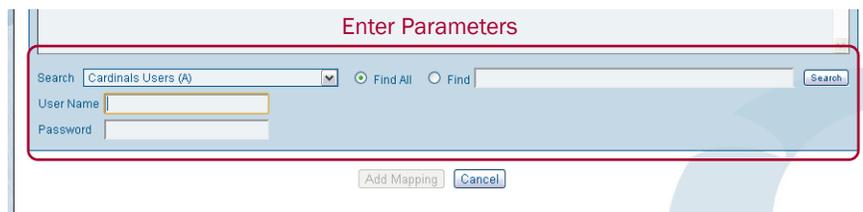
2. Select a mapped **user/group** to be changed, and click **Change Mapping**.



3. At the **Change Mapping** page, to find the user/group you want to map to, select the Windows domain **user or group list**, the search scope, enter a search string if needed, and click **Search**.
 - To search for a specific user or group, use either **Find All** or a **Find** search string (wildcard “*” before or after is allowed).

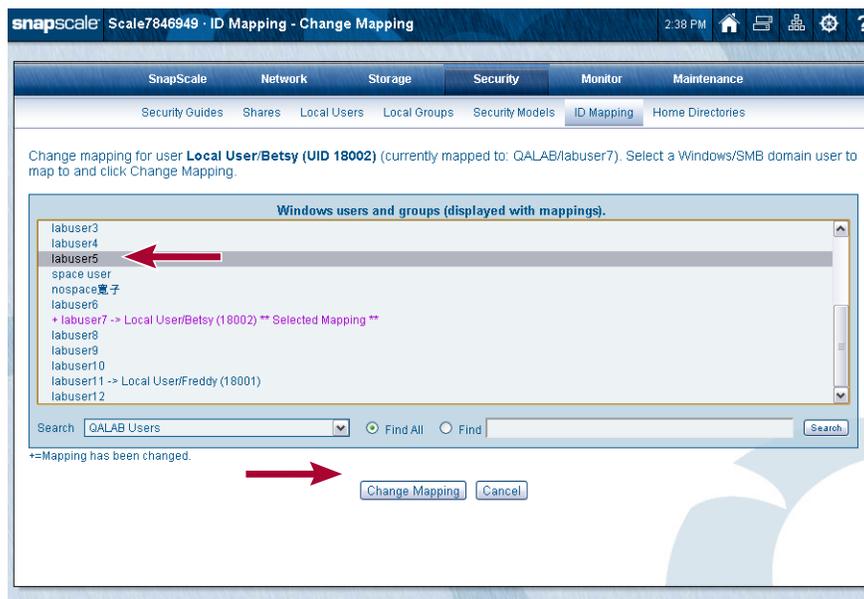


- For domains that REQUIRE authentication (showing an “(A)” after the name), select the domain name, enter the user name and password for that domain, and use either **Find All** or a **Find** search string (using the first few letters of the user/group name).



- On the rare occasion you need to search for a Windows domain that's not listed (“remote domain”), select a Windows domain from the **Search** drop-down list through which to search, then enter in the **Find** box the name of the remote domain, followed by a slash (/) or backslash (\) and the user name for which you are searching (for example, **remote_domain\user_name**). After you click **Search**, another authentication prompt may be presented to authenticate with the remote domain.

- From the search results, select the Windows domain **user/group** to which you want to re-map the local or NIS user, and click **Change Mapping**.



- Repeat [Steps 1–4](#) until all **changes** are made.
- Click **OK** to save changes (or **Cancel** to reset).
- When done with all your mappings, click **OK** to activate them.
- At the confirmation page, click **Save Changes**.
- At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

See [Filesystem Updates](#) on [page 6-37](#) for more details.

Auto Mapping

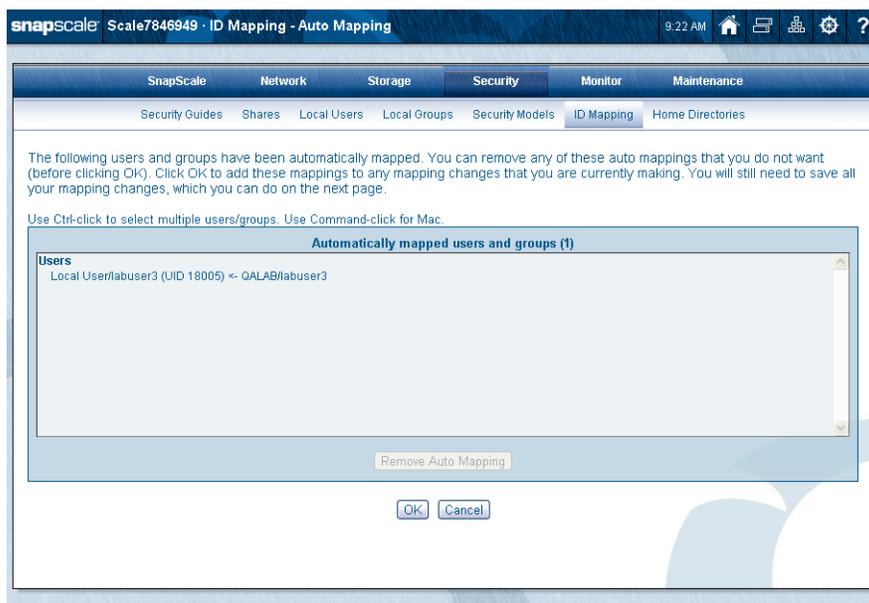
Auto mapping generates a list of ID mappings for Windows users and groups that have the same name as your local or NIS users and groups (local has precedence over NIS).

- Click **Auto Mapping** to generate a **list** of Windows domain users/groups that have the same name as your local or NIS users and groups.
Domain, local, and NIS user/group lists are compared. The matches are automatically queued. Users and groups already mapped are not affected.

- At the **Auto Mapping** confirmation page, click **View Auto Mappings** to continue.



- At the summary page, verify the **mappings** and remove (**Remove Auto Mapping**) any users or groups you do not want to map.



- Click **OK** to save changes (or **Cancel** to reset).
- When done with all your mappings, click **OK** to activate them.
- At the confirmation page, click **Save Changes**.
- At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

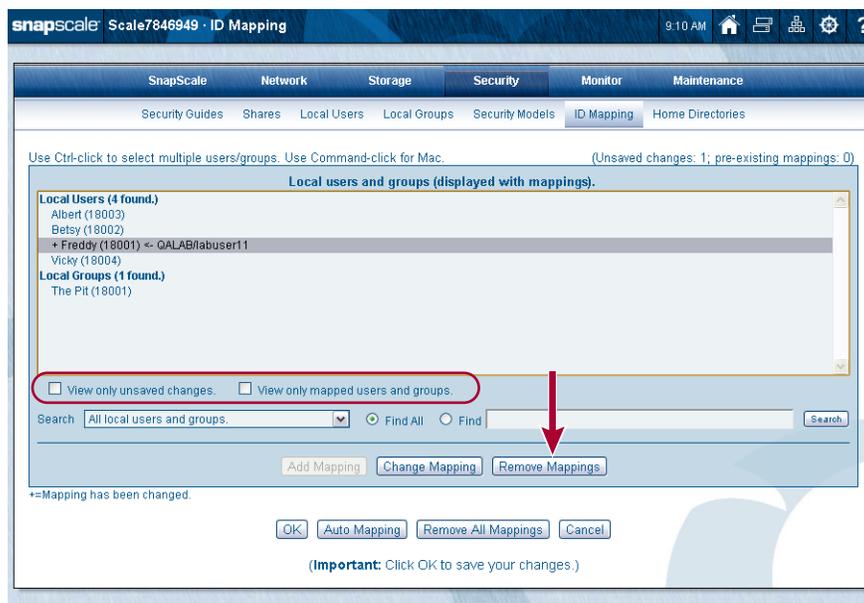
See [Filesystem Updates](#) on [page 6-37](#) for more details.

Remove Mappings

User mappings can be removed individually or all at once. Once removed, they can not be restored but must be added back using [Add Mapping](#) on [page 6-26](#). You also have the option to update the filesystem after removing the ID mappings.

Remove Individual Mappings

1. At the default **ID Mapping** page, select one or more **users/groups** you wish to unmap. To make it easier to find mappings for removal, check **View only mapped users and groups** to display only local or NIS users or groups that have been mapped.
2. Click **Remove Mappings**.



- At the confirmation page, verify the **users/groups** listed and click **Remove Mappings**.



The selected mappings are removed and the default page is displayed with the users/groups that were unmapped in purple with a plus (+) in front of their name.

- Click **OK** to save changes (or **Cancel** to reset).
- When done with all your mappings, click **OK** to activate them.
- At the confirmation page, click **Save Changes**.
- At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

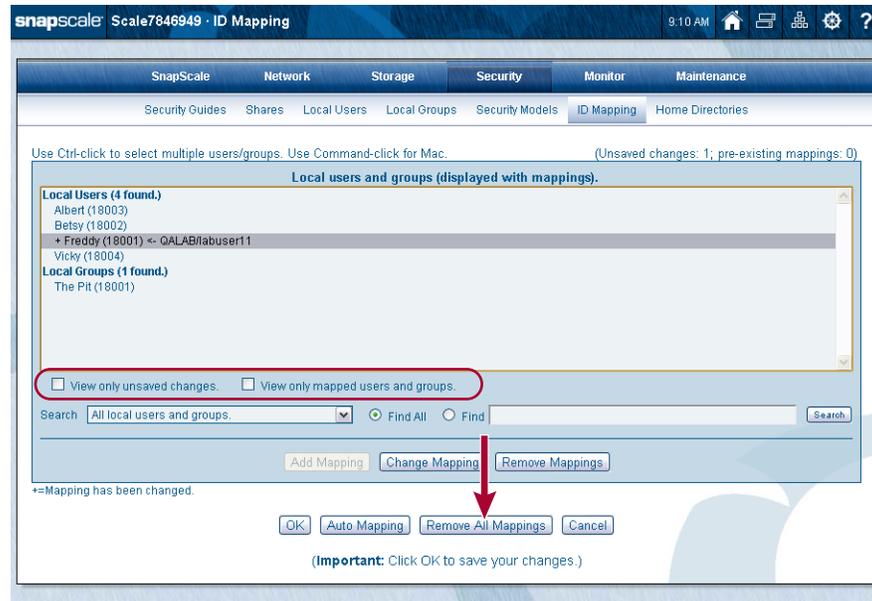
 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

See [Filesystem Updates](#) on [page 6-37](#) for more details.

Remove All Mappings

The **Remove All Mappings** button allows you to remove all ID mappings on the cluster. Click this only if you want to remove all ID mappings. If there are no mappings, the button is grayed out.

- At the default **ID Mapping** page, click the **Remove All Mappings** button.
If needed, check **View only unsaved changes** to display only mapping changes that have not yet been applied. Check **View only mapped users and groups** to display only local or NIS users/groups that have been mapped to a Windows domain user or group.



- At the confirmation page, click **Remove Mappings**.



All the mappings are removed and the default page is displayed with the users/groups that were unmapped in purple with a plus (+) in front of the names.

- Click **OK** to save changes (or **Cancel** to reset).
- When done with all your mappings, click **OK** to activate them.
- At the confirmation page, click **Save Changes**.
- At the filesystem update option page, choose either **Update Filesystem** or **Do Not Update Filesystem**.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

See [Filesystem Updates](#) on page 6-37 for more details.

Remove Missing ID Mappings

If the cluster has mappings for users or groups that no longer exist, the following warning message may be displayed at the top of the main **ID Mappings** page:



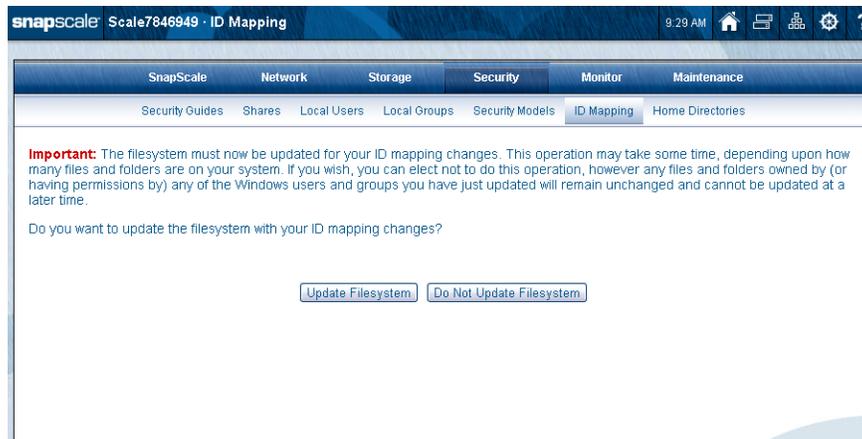
1. Click the **Click here** link in the warning message to display the **Remove Missing Mappings** page.
2. Click **Remove Missing Mappings** to clear the missing mappings from the system. A confirmation is shown on the **ID Mapping** main page.



3. Click **OK** to save changes.

Filesystem Updates

After making any changes to ID mappings, you are presented with a filesystem update option page, where you can choose either **Update Filesystem** or **Do Not Update Filesystem** options.

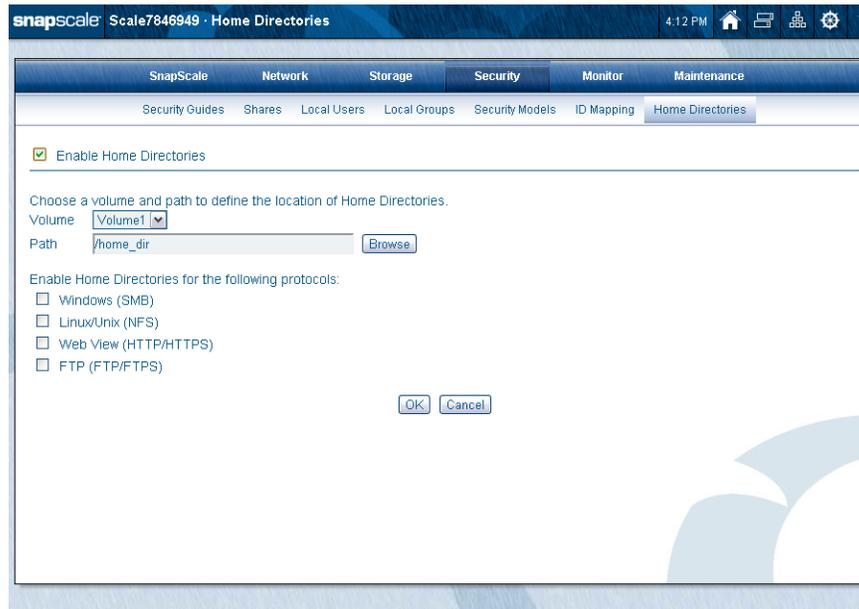


If you choose **Update Filesystem**, UID and GID ownership on files and SIDs in ACLs are updated to reflect the ID mapping operation.

 **IMPORTANT:** Updating may take some time, depending upon how many files and folders are on your system. If you elect not to do this operation, any files and folders owned by (or having permissions by) any of the Windows users and groups you have just updated will remain unchanged and cannot be updated at a later time.

Home Directories

To enable Home Directories, go to **Security > Home Directories** and check **Enable Home Directories**. Choose the volume, path, and protocols you want.



The Home Directories feature creates a private directory for every local or Windows domain user that accesses the system. When enabling Home Directories (from **Security > Home Directories**), the administrator creates or selects a directory to serve as the home directory root. When a user logs in to the cluster for the first time after the administrator has enabled Home Directories, a new directory named after the user is automatically created inside the home directory root, and is configured to be accessible only to the specific user and the administrator.

Depending on the protocol, home directories are accessed by users either via a user-specific share, or via a common share pointing to the home directory root.

Home directories are supported for SMB, NFS, HTTP/HTTPS, and FTP/FTPS. They are accessed by clients in the following manner:

- For SMB, HTTP/HTTPS and FTP/FTPS, users are presented with a virtual share named after the user name. The virtual share is visible and accessible only to the user. Users are not limited only to their virtual shares; all other shares on the cluster continue to be accessible in the usual fashion.
- For NFS, the home directory is exported. When a user mounts the home directory root, all home directories are visible inside the root, but the user's home directory is accessible only by the user and the administrator.

NOTE: If desired, UNIX clients can be configured to use a Snap Home Directory as the local user's system home directory. Configure the client to mount the home directory root for all users, and then configure each user account on the client to use the user-specific directory on the SnapScale as the user's home directory.

If ID Mapping is enabled, domain users and local users mapped to the same user are directed to the domain user's home directory. In some cases, data in the local user's home directory is copied to the domain user's home directory:

- If a local user home directory accumulates files before the local and domain users are mapped, and if the domain user's home directory is empty, the local user's files are copied to the domain user's home directory the first time the local user connects after the users are mapped.
- If both the local and domain user home directories accumulate files before the local and domain users are mapped, the files in the local user's home directory are not copied to the domain user's home directory.

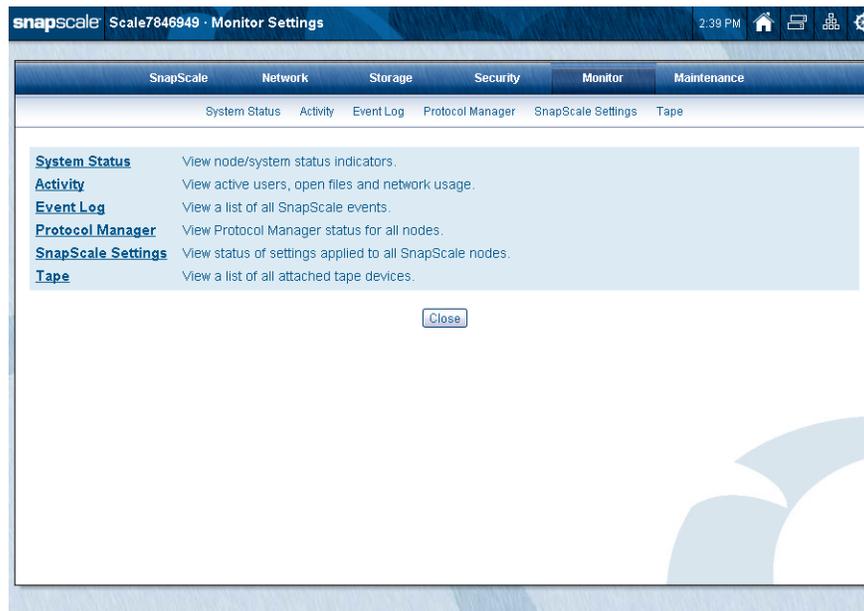
Configure Home Directories

Check or complete the following fields, and click **OK**.

Field	Description
Enable Home Directories	Check to enable Home Directories for local users and activate the options. Remove the check to disable.
Volume	Select the volume where the Home Directories will be located. NOTE: Be sure the volume you select has enough disk space. Once Home Directories are placed, they cannot be moved.
Path	Provide the path to the Home Directories or click Browse to create a new folder. The default path is <code>/home_dir/</code> .
Protocols	Check each of the protocols where Home Directories will be enabled.

NOTE: Do not put Home Directories on a volume that might be deleted. If you delete the volume, you will also delete the Home Directories.

This chapter addresses the options for monitoring the SnapScale cluster. Here you can view the system status and other activities.



Topics in System Monitoring:

- [System Status](#)
- [Activity](#) submenus:
 - [Active Users](#)
 - [Open Files](#)
 - [Network Monitor](#)
- [Event Log](#)
- [Protocol Manager](#)
- [SnapScale Settings](#)
- [Tape](#)

System Status

Use the **System Status** page (**Monitor > System Status**) to assess the hardware status and key information of the cluster member nodes.



SnapScale Status

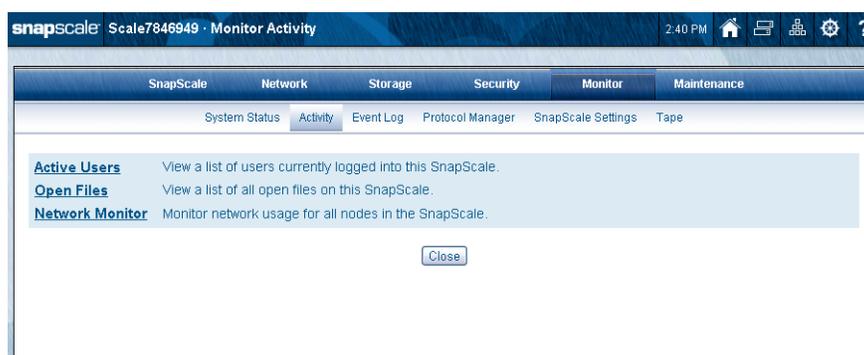
The following status fields are displayed for each node that is part of the SnapScale cluster. Any critical messages are displayed in a **red** font.

Field	Description
Node Name	Name of the node: Node <code>nnnnnn</code> (where <code>nnnnnn</code> is your node number). Example: Node2302216.
Node Model	Node hardware model.
OS Version	The version of RAINcloudOS currently loaded on the node.
Hardware	The node's hardware platform ID.
Node Number	Number derived from the MAC address of <i>Ethernet 1</i> port that is used as part of the node name.
BIOS	The BIOS version for the node.
Serial Number	Unique number assigned to the node.
JVM	The Java Virtual Machine version.
Uptime	The amount of time the node has been up (since the last reboot) in "days:hours:minutes" format.
Memory	Amount of system RAM.
CPU	The type of central processing unit for the node's first CPU.
Client Network	Details on the node's client Ethernet connections.
Storage Network	Details on the node's storage Ethernet connections.

Field	Description
Ambient Temp.	The temperature of the space inside the chassis.
CPU Temp.	Current CPU temperature.
Power Supply	The status of power supply modules
Fan Status	The status of fan modules.

Activity

The **Activity** tab provides access to a submenu of options and features for monitoring activity on the cluster.



This submenu is used to access three other pages—**Active Users**, **Open Files**, and **Network Monitor**.

Active Users

This option is used to view read-only details on the active users logged on to each of the nodes on the cluster.



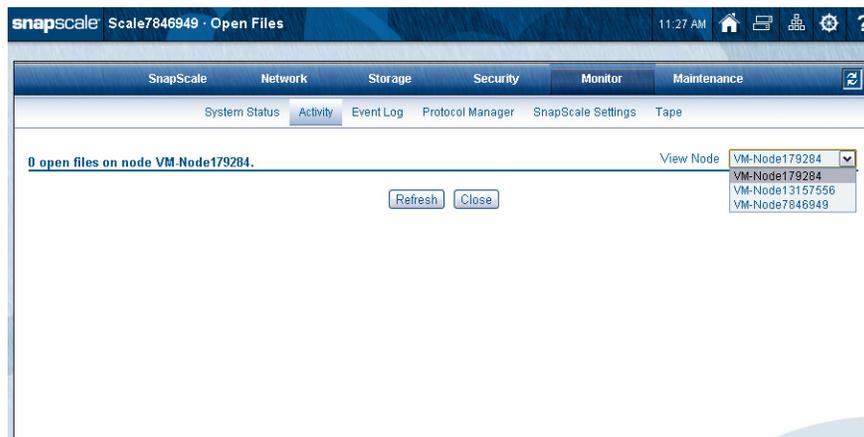
Information available on this page includes user names of all active users, their workstation names, authorization, the number of open files they have on the node, the protocol, and when they logged on. Columns can be sorted in ascending or descending order by clicking the column head.

Use the drop-down list on the upper right to select whether to view all the nodes or individual nodes. Close the page to return to the **Activity** tab.

NOTE: Active users are not displayed for HTTP or NFS.

Open Files

Use this option to view read-only details on the open files on a specific node.



Use the drop-down list on the upper right to choose a different node to view. Close the page to return to the **Activity** tab.

Network Monitor

This feature can be used to monitor the network utilization on both the Client and Storage networks. The default is disabled to minimize system overhead.



Go to **Monitor > Activity > Network Monitor** and click **Enable Network Monitoring** to activate.

Use the **View** drop-down list on the upper right corner of the table to display a chart of current usage, average usage, or utilization through time for one of several time periods in the past hour. The data is refreshed every 11 seconds.

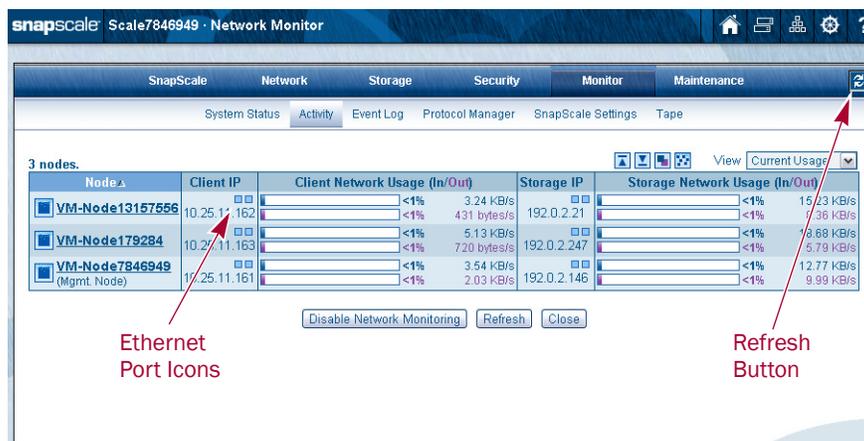
A tool tip message is displayed when you mouse-over the usage bars for a specific node. It shows information about the bonded interface including the maximum bond speed and type, total throughput, and (if Average view is selected) the time that averaging was started.

NOTE: All throughput values are displayed in bytes (and not bits) per second.

To stop the monitoring, click **Disable Network Monitoring** and confirm. Close the page to return to the **Activity** submenu.

Current Usage View

The **Current Usage** view on the **Network Monitor** page shows the input and output (In/Out) usage for both the Client and Storage networks for all the nodes.



While the information is automatically refreshed at regular intervals, you can use either the **Refresh** button at the bottom or the Refresh icon (🔄) on the right corner of the tab bar to manually refresh the information.

Position the cursor over the dual Ethernet port icons (□□) displayed for each node's Client and Storage network bond to view the individual port status. An Ethernet port icon is displayed in yellow if the port has no link.

To the left of the View drop-down menu are icons used to configure the display:

Icon	Name	Description
	High Water (HW) Marks	Click icon to turn on the high water mark option that shows the peak usage. The icon turns green and a lighter colored bar in each of the In and Out bars shows the maximum usages. A high-water mark specific tool-tip message is displayed when you mouse-over the high-water mark portion of the usage bars for a specific node. It shows information about the high-water mark's throughput as well as the date and time when it occurred. A message is displayed above the table to the right of the node count stating the time span for the HW mark.
	High Water Reset	Click icon and confirm to reset the high water marks for the usage. This icon is grayed out when HW marks are turned off.
	Individual/Total Network Usage	When blue, the numbers reflect the individual input and output usage. Click to show the combined total usage (the icon turns green). Click again to return to the individual readings.
	Cluster Total Usage	Click icon to show usage for the entire SnapScale cluster (combining all the node usage amounts). The icon turns green and a new row for the cluster is shown at the top of the table. Click again to remove the row.

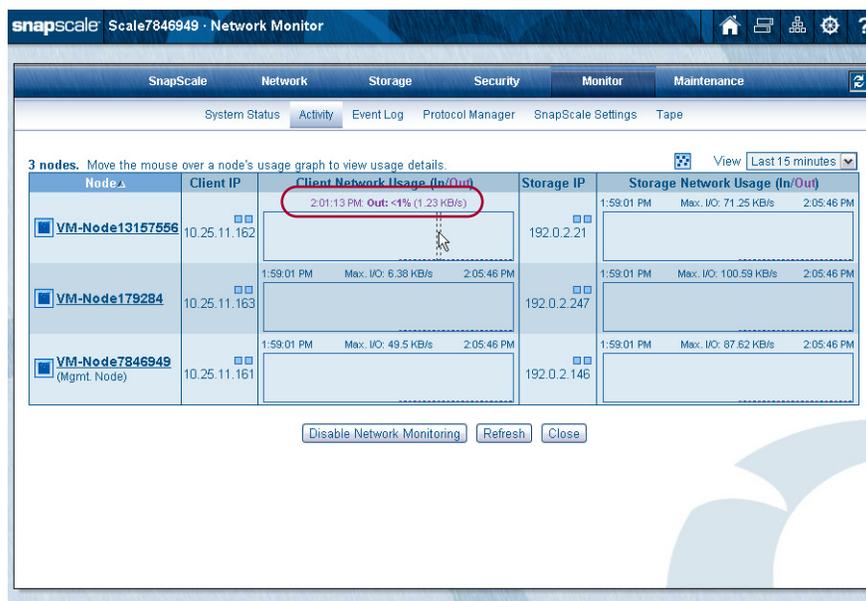
Average Usage View

This view is similar to Current View but shows the average usage in the usage columns. It also adds an icon (🔄) to reset the average usage numbers.

Specific Time Views

To view historical usage over a specified time period in the last hour, select the time period from the **View** drop-down menu. The options are 5, 10, 15, 30, 45, or 60 minutes. The start time, maximum I/O, and end time are displayed about the graphs. Each graph can show up to 26 (averaged) samples of I/O data for the selected time period.

Position the cursor over a usage graph to view usage details at a specific time in that range.



Click the icon above the table to add the network activity for the entire SnapScale cluster.

Event Log

Use the **Event Log** page to view a log of operations performed on the cluster.

Change the following fields and click Refresh to specify how the log is displayed.

View Log Severity Display Last Days Most Recent First

Severity Legend: E=Error, W=Warning, I=Information

S	Time	Message	Source
I	09/26 3:59:30 PM	Node2414538: root prepared the cluster FS - 70% complete (04:18.279 elapsed)	EventSystem
I	09/26 3:59:30 PM	Node2414538: root confirmed cluster fs mounted on all nodes - 60% complete (04:17.872 elapsed)	EventSystem
I	09/26 3:58:57 PM	Node2414528: the cluster FS is mounted locally - 50% complete (03:40.835 elapsed)	EventSystem
I	09/26 3:58:23 PM	Node2414528: peersets all report running - 40% complete (03:40.684 elapsed)	EventSystem
I	09/26 3:58:22 PM	Node2414532: the cluster FS is mounted locally - 50% complete (02:40.844 elapsed)	EventSystem
I	09/26 3:58:02 PM	Node2414532: peersets all report running - 40% complete (02:40.704 elapsed)	EventSystem
I	09/26 3:59:30 PM	Node2414528: 12 volumes registered in cluster table - 30% complete (03:10.266 elapsed)	EventSystem
I	09/26 3:59:30 PM	Node2414532: 12 volumes registered in cluster table - 30% complete (02:10.275 elapsed)	EventSystem
I	09/26 3:58:57 PM	Node2414538: the cluster FS is mounted locally - 50% complete (03:32.849 elapsed)	EventSystem
W	09/26 3:58:23 PM	Node2414538: Set cluster to read-write	EventSystem
I	09/26 3:59:30 PM	Node2414538: peersets all report running - 40% complete (03:32.501 elapsed)	EventSystem
I	09/26 3:59:30 PM	Node2414538: 12 volumes registered in cluster table - 30% complete (03:02.029 elapsed)	EventSystem
W	09/26 3:58:57 PM	Node2414528: Received SYSTEM_SHUTDOWN. Shutting down all services.	EventNode
W	09/26 3:58:23 PM	Node2414538: Received SYSTEM_SHUTDOWN. Shutting down all services.	EventNode
I	09/26 3:58:22 PM	Node2414538: Cluster has started	EventSystem
I	09/26 3:58:02 PM	Node2414538: cluster is formed!! - 90% complete (05:19.166 elapsed)	EventSystem

Entries are color coded according to severity as described in the following table:

Color	Icon	Entry Type
Red		Error (E)
Yellow		Warning (W)
(no color)		Informational or Unclassified (I)

Filter the Log

Edit the following fields as appropriate, then click **Refresh**.

Option	Description
View Log	Select to view either the SnapScale cluster-wide or node-specific logs. The SnapScale option shows general cluster-related log messages while the node-specific options show log messages specific to the selected node.
Severity	Select the type of alerts and information you want to view.
Display Last __ Days	Enter the number of days' worth of entries you want to view.
Most Recent First	Check this box to start the list with the most recent entry; deselect to start the list with the oldest entry.

Protocol Manager

Protocol Manager manages networking protocols and IP address assignment across the entire SnapScale. If a node fails or is removed from the SnapScale, Protocol Manager handles automatic IP address reassignment to maintain client access to data.

The screenshot shows the SnapScale Protocol Manager interface. At the top, there's a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Below that, there's a sub-navigation bar with System Status, Activity, Event Log, Protocol Manager (selected), SnapScale Settings, and Tape. The main content area displays a message: "Protocol Manager manages networking protocols and IP address assignment across the entire SnapScale. If a node fails or is removed from the SnapScale, Protocol Manager handles automatic IP address reassignment to maintain client access to data." Below this message, it says "3 nodes." and shows a table with the following data:

Node	Description	Status	IP Address
VM-Node13157556	-	OK	10.25.11.162
VM-Node179284 (Mgmt. Node)	-	OK	10.25.11.163
VM-Node7846949 (IP Manager)	-	Banned, Disabled, Inactive	10.25.11.161

Below the table, there is a note: "*Note: An online node's status may be unhealthy, or the node's IP address may be unassigned, temporarily for several minutes after the SnapScale is first created, a new node is added, or the SnapScale is restarted. If an online node's status is displayed as unhealthy (any status other than *OK*) for longer than 30 minutes, please contact Overland Support." At the bottom of the interface, there are "Refresh" and "Close" buttons.

The following table addresses the possible status:

Status	Description
OK	This node is fully functional.
Disconnected	This node could not be connected through the Storage network and is currently not participating in the cluster. If there is a public IP address associated with this node it should have been taken over by a different node. No services are running on this node.
Banned	This node failed too many recovery attempts and has been banned from participating in the cluster temporarily. Any public IP addresses have been taken over by other nodes.
Disabled	This node has been administratively disabled. This node is still functional and participates in the cluster but its IP addresses have been taken over by a different node and no services are currently being hosted.
Unhealthy	A service provided by this node is malfunctioning. The node itself is operational and participates in the cluster, however its public IP addresses have been taken over by a different node and no services are currently being hosted.
Stopped	A node that is stopped does not host any public IP addresses, and does not participate in the cluster.
PartiallyOnline	A node that is partially online participates in the cluster like a node that is OK. Some network interfaces which serve public IP addresses are down, but at least one interface is up.

SnapScale Settings

When cluster settings are configured in the Web Management Interface, success or failure of the operation is determined by the attempt to perform it on the Management node. If successful, the same configuration operation is pushed to all member nodes in the background.

The **SnapScale Settings** page displays a list of settings that have been applied to the nodes in the cluster and the status of each setting. When you make changes to your SnapScale via the Web Management Interface, the settings are applied to the Management node first to determine success or failure of the configuration, then the settings are applied to the other nodes in the background. When a SnapScale setting has not been applied yet, its status is displayed as **Pending**. When a SnapScale setting fails to be applied, its status is displayed in detail and the failed settings are automatically re-applied until they are successful.

The initial view is compressed to show all nodes and a count of the settings:

The screenshot shows the SnapScale Settings page with the following data:

Node	Settings	Status	Time
VM-Node13157556	(8)	Settings successfully applied.	2013-09-06 2:30:47 PM
VM-Node179284 (Mgmt. Node)	(8)	Settings successfully applied.	2013-09-06 2:30:11 PM
VM-Node7846949	(8)	Settings successfully applied.	2013-09-06 2:30:46 PM

Buttons: Refresh, Close

Click the upper right text that says **View is: Compressed** to switch to the expanded view:

The screenshot shows the SnapScale Settings page in expanded view. The page title is "Scale7846949 · SnapScale Settings" and the time is 2:45 PM. The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The main content area shows "SnapScale Settings" and a table of settings for 3 nodes. The table has columns for Node, Settings, Status, and Time. The settings are grouped by node, and the status for all is "Settings successfully applied." The time of application is listed for each setting.

Node	Settings	Status	Time
VM-Node13157556	Users & Groups	Settings successfully applied.	2013-09-06 2:30:47 PM
	Snapshot Schedules	Settings successfully applied.	2013-09-06 12:03:45 PM
	Snap EDR	Settings successfully applied.	2013-09-04 5:33:43 PM
	Share Access	Settings successfully applied.	2013-09-04 5:32:23 PM
	NFS Exports	Settings successfully applied.	2013-09-04 5:32:22 PM
	Shares	Settings successfully applied.	2013-09-04 5:32:22 PM
	Server	Settings successfully applied.	2013-09-04 5:32:21 PM
VM-Node179284 (Mgmt. Node)	Volumes	Settings successfully applied.	2013-09-04 5:32:14 PM
	Users & Groups	Settings successfully applied.	2013-09-06 2:30:11 PM
	Snapshot Schedules	Settings successfully applied.	2013-09-06 12:03:45 PM
	Snap EDR	Settings successfully applied.	2013-09-04 5:33:43 PM
	Server	Settings successfully applied.	2013-09-04 5:32:22 PM
	Share Access	Settings successfully applied.	2013-09-04 5:32:21 PM
	NFS Exports	Settings successfully applied.	2013-09-04 5:32:21 PM

Buttons: Refresh, Close

Use the scroll bar on the right side to view all the data. In expanded mode, detailed information on each node is available:

- Each line reports the setting category, status, and date/time the setting was applied.
- If an operation is still in progress on a node, the **Status** will be set to **Pending** with a yellow background.
- If an operation failed on a node, the **Status** will have an error message and a red background.
- Column headers can be clicked to sort by **Node**, **Settings**, **Status**, or **Time**.
- When sorting by **Status** or **Node**, settings are grouped by node.

Click the **View is: Expanded** text to revert back to the compressed view. The displayed view from that point on will be the last view selected. Clicking the column heading resorts the table on that function.

Tape

Use the **Tape Monitor** page (**Monitor > Tape**) to view read-only details on the SCSI and USB tape devices attached to each node.

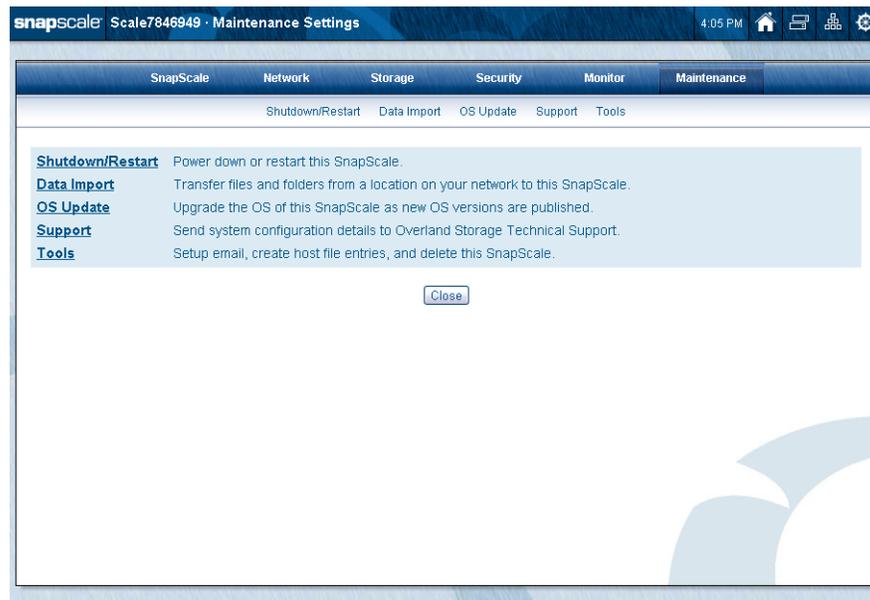


Use the drop-down list on the upper right to select whether to view all the nodes or individual nodes. Close the page to return to the **Monitor** page.

The following table details the fields:

Field	Description
Device Model	The manufacturer's model for the device.
Device Type	Type of tape device: either Sequential-Access (tape drive) or Medium-Changer (for example, robotic arm for a tape library).
Device Name	Name of the device node to which the device is bound.
Connection	Identifies the connection type: SCSI or USB.
Bus	Bus number indicating which physical interface (for example, SCSI card) the device is connected to.
ID	ID number (SCSI only)
LUN	LUN identifier (SCSI only)

Clicking the **Maintenance** tab on the Web Management Interface displays options used to maintain this SnapScale cluster. There is also a **Tools** submenu of special, related options.

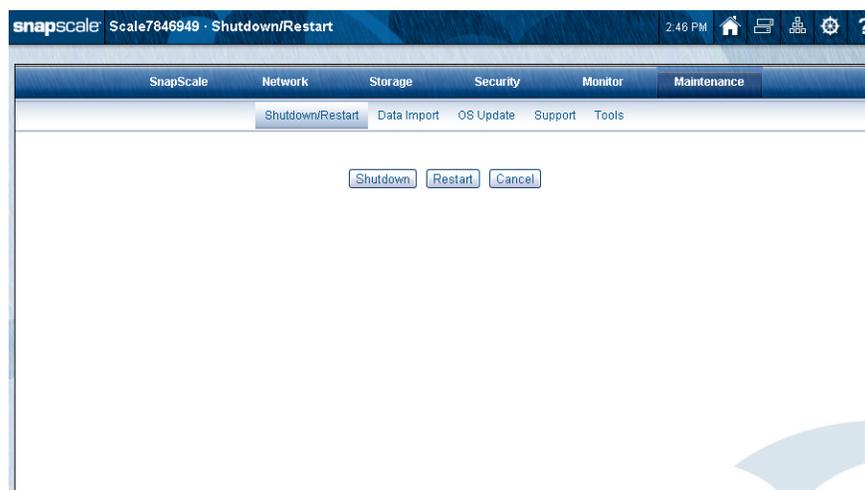


Topics in Web Management Interface

- [Shutdown and Restart](#)
- [Data Import](#)
- [OS Update](#)
- [Support](#)
- [Tools](#) submenus:
 - [Email Notification](#)
 - [Host File Editor](#)
 - [Delete SnapScale Cluster](#)

Shutdown and Restart

Use the **Shutdown/Restart** page to reboot or shut down the cluster.



Click one of the following buttons:

- **Shutdown** – Shuts down and powers off all nodes in the cluster.
- **Restart** – Reboots the cluster via a controlled shutdown and restart.

Manually Powering Nodes On and Off

 **CAUTION:** To prevent possible data corruption or loss, it is NOT recommended to directly power down any nodes that are part of a SnapScale cluster. When powering down a cluster, always use the Shutdown button that can be found under Maintenance > Shutdown/Restart in the Web Management Interface.

Use the power button on the front to power ON and power OFF (in an emergency) a node:

- To turn the node ON, press the power button on the front of the node.
The node takes a few minutes to initialize. A solid green system/status LED indicates when the system is up and running.
- To turn the node OFF, press and release the power button to begin the shutdown process. Do not depress this button for more than four seconds.

NOTE: All SnapScale nodes have a persistent power state. When a physical loss of power occurs, the node returns to the same operation being performed when the power went out. Therefore, if the node is powered down prior to a power loss, it will remain powered down when the power is restored, and if it was powered on prior to a power loss, it will power back on when power is restored.

Data Import

Use the **Data Import** page (**Maintenance > Data Import**) to import (migrate) data from another SnapScale cluster, SnapServer, or other computer that supports CIFS or NFS (v2 or v3) to this cluster.

Windows/SMB Page:

The screenshot shows the SnapScale Data Import interface for the Windows/SMB protocol. The page title is "Scale7846949 · Data Import" and the time is 2:47 PM. The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The "Data Import" tab is active, with sub-tabs for Shutdown/Restart, OS Update, Support, and Tools. The main content area contains the following fields and options:

Use Data Import to copy or move files and folders from a location on your network (Source) to this SnapScale (Target).

Source:

Network Protocol: (specifies how to communicate with host)

Auth. Name:

Auth. Password:

Host:

Share:

Path:

Target (This SnapScale):

Volume:

Path:

Options:

Import Type:

Include all sub-folders (if source path specifies a folder)

Overwrite existing target files and folders (that have identical names as the source files and folders)

Preserve file/folder permissions

Verify imported data (takes twice as long)

Note: You can setup [Email Notification \(administrative operation event\)](#) to be notified when a Data Import operation is complete.

NFS Page:

The screenshot shows the SnapScale Data Import interface for the NFS protocol. The page title is "Scale7846949 · Data Import" and the time is 10:26 AM. The navigation menu includes SnapScale, Network, Storage, Security, Monitor, and Maintenance. The "Data Import" tab is active, with sub-tabs for Shutdown/Restart, OS Update, Support, and Tools. The main content area contains the following fields and options:

Use Data Import to copy or move files and folders from a location on your network (Source) to this SnapScale (Target).

Source:

Network Protocol: (specifies how to communicate with host)

User Name: (Snap local or NIS user)

Host:

Export:

Path:

Target (This SnapScale):

Volume:

Path:

Options:

Import Type:

Include all sub-folders (if source path specifies a folder)

Overwrite existing target files and folders (that have identical names as the source files and folders)

Preserve file/folder permissions

Verify imported data (takes twice as long)

Note: You can setup [Email Notification \(administrative operation event\)](#) to be notified when a Data Import operation is complete.

If an error is encountered during the import (for example, a file or folder is locked and cannot be imported), the utility records the error in a log, and continues the operation. When the import is completed, the administrator can view the log of import errors. Once the errors have been corrected, the administrator can return to the main page, and recreate the import. With the exception of the password, all fields are still be populated with the specifications of the last job.

The following data import options can be specified:

- Import type: copy or move data
- Include subfolders
- Overwrite existing files
- Preserve the original permissions settings

NOTE: If you elect to preserve original permissions settings, review Preserving Permissions on page 8-6.

- Verify imported data

NOTE: If you elect to verify imported data, all data is read twice, once for import and once for comparison to the copied data. This could be a lengthy process.

Setting Up a Data Import Job

Before setting up a data import job, be sure to specify a user identity for the operation that has full access to all files on the source, regardless of permissions set:

- For Windows import, specify an administrator or member of the Windows server/domain administrators group.
- For NFSv2/3 import, consider using the user root, and configuring the NFS export on the source to `no_root_squash` for the IP Address of the node for the duration of the import.

NOTE: Only one import job can run at a time.

To create a data import job, perform the following procedure:

1. On the **Data Import** page, complete the required **information** for both the source and target.

Option	Description
Source:	
Network Protocol	<p>Protocol that the cluster uses to connect to the source server. Use the drop-down list to select:</p> <p>NOTE: If you are importing via SMB, SMB must also be enabled on the target cluster (enable at Network > Windows/SMB).</p> <ul style="list-style-type: none"> • Windows (SMB) – Uses SMB for Windows with the source data on a Windows root directory. Default option. • NFS – NFSv2/3 for UNIX/Linux-based servers or RAINcloudOS nodes with source data on a UNIX root directory.
Auth. Name & Auth. Password / User Name	<ul style="list-style-type: none"> • For the Windows (SMB) network protocol, enter both the Auth. Name and Auth. Password (Windows user name and password to log in to the server over SMB). • For the NFS network protocol, enter the User Name (node local user name or NIS user, representing the UID used to perform the operation over NFS).
Host	Enter the name or IP address of the source computer you are importing data from.
Share/Export	<p>Specify the share (Windows) or export (NFS) on the source server containing the data you want to import.</p> <p>NOTE: Wildcards are not supported when specifying the source share to import.</p>

Option	Description
Path	Enter the path to the file or folder you want to import. If you are importing the entire share, you can leave the Path field blank. NOTE: Wildcards are not supported when specifying the path to import.
Target (The SnapScale):	
Volume	Specify the cluster volume where you want the data imported.
Path	Specify the path to the directory where you want the data imported.
Options:	
Import Type	Options for the import data are to Copy (source data is maintained) or Move (source data is removed during copy). If Verify Imported Data is enabled, the Move option removes the original data after the verification is complete. The default is Copy. NOTE: If you select to Move rather than Copy data, it is strongly recommended that you also select the Verify Imported Data option.
Include All Sub-folders	If the folder you select for import contains subfolders, selecting this option imports all files and folders underneath this folder (checked by default). If disabled, <i>only</i> the files in this folder are imported.
Overwrite Existing Target Files & Folders	If any files/folders on the target have identical names with files/folders on the source, checking this option overwrites those files/folders during import (checked by default.)
Preserve File/Folder Permissions	Selecting this option retains the source permissions when the files/folders are imported to the target (unchecked by default). NOTE: Before selecting this option, review Preserving Permissions on page 8-6.
Verify Imported Data	Selecting this option causes all source data to be read twice, once to write to the target and once to perform a binary comparison with the data written (unchecked by default). If enabled, and if the Import Type is Move , files on the source are only removed after verification. Otherwise, files are removed during the process of copying them to the target. If you select to move files rather than copy them, it is strongly recommended that you enable the Verify Imported Data option. If a file mismatch occurs during verification, the target file is moved to a <code>data_import_verify_failures</code> directory on the root of the same volume. Check the failed file to determine the problem, then run the import again with Overwrite Existing Target Files & Folders deselected (so you do not re-copy files that have already been copied and verified). NOTE: Depending upon how much data is being imported, verifying imported data can be a lengthy process.
Email Notification	Clicking the email notification link takes you to the Email Notification page (for more information, see Email Notification on page 8-14). Fill in notification information and check the box next to Administrative Operation Event in order to receive an email when the import operation is complete.

2. Once you have completed the import information, click the **Start Import** button to begin the import. You can see the progress of the import, the estimated time until completion, and the import log on the secondary **Data Import** page.
3. When the import is complete, click the **View Log** button to see details of all errors. Click the **Data Import Error Log** link to download the entire log.

Stopping an Import Job

To stop the import at any time, click the **Stop Import** button on the **Data Import** secondary page. If a file was in the process of being copied, the partially-copied file on the target is removed.

Recreating an Import Job

The **Data Import** log records all errors that occurred during import. You can import just the files and folders that were not imported during the original job due to an error condition (for example, the file was locked).

1. Review the **Data Import errors log** and correct all error conditions (such as unlock a file).
2. Reopen the **Data Import** page. All fields (except the password) from the last import will still be visible on the page.

By default, all files will be re-imported. If you want only to import those files that failed to import the first time, you can disable the **Overwrite Existing Target Files** option. However, make sure that all problematic files from the first import are deleted from the target so they can be re-imported.

NOTE: If an import failed, it is strongly recommended that you enable the [Verify imported data option for the re-importation](#).

3. Enter your password and click **Start Import** to run the import again.

Preserving Permissions

The types of permissions retained will differ, depending on which import scenario is applied.

Importing from a Windows Security Model to a Windows Personality Directory

If you are importing from a Windows server (or other type of server that follows the Windows security model) to a Windows personality directory, permissions are retained exactly as they exist on the source. However, as is the case when moving files with permissions between Windows servers, permissions for users who are unknown on the target are retained but not enforced. This includes permissions for:

- Local users on the source machine.
- Domain users for domains unknown to the cluster (for example, trusted domains, if the cluster is not configured to support trusted domains).
- Certain built-in Windows users and groups.

Importing from a UNIX Security Model to a UNIX Personality directory

If you are importing from a UNIX server to a UNIX personality directory, UNIX permissions for UIDs/GIDs are copied exactly from source to target; thus, identities of the users and groups are best retained if the SnapScale cluster belongs to the same NIS domain as the UNIX server.

Importing Between Conflicting Security Models

When importing from a UNIX source to a Windows security model target, UNIX permissions are retained and the security personality on the resulting files and directories will be UNIX.

However, when importing from a Windows source to a UNIX security model target, permissions cannot be retained (since UNIX root directories are required to be UNIX personality throughout). Files and directories will inherit the UNIX personality and will have a set of default UNIX permissions.

Importing from a SnapServer or SnapScale Cluster

When importing from a SnapServer or another SnapScale cluster, it is recommended that you maintain the same security model on the target that you have on the source.

- If your source uses a Windows security model and has permissions assigned to Windows domain users, use a Windows (SMB) connection for import. Windows permissions are retained exactly as they are on the source, with the same enforcement limitations for unknown users as for importing from Windows servers (see [Importing from a Windows Security Model to a Windows Personality Directory](#) on page 8-6).
- If your source server or cluster uses a UNIX security model and has permissions assigned to local or NIS users, use an NFS connection for import.

NOTE: Local users who have UNIX permissions on the source are not created on the target with the same UIDs.

OS Update

Use this page to install updates to RAINcloudOS and other installed software, and to configure RAINcloudOS to automatically check for updates to RAINcloudOS and Snap EDR.

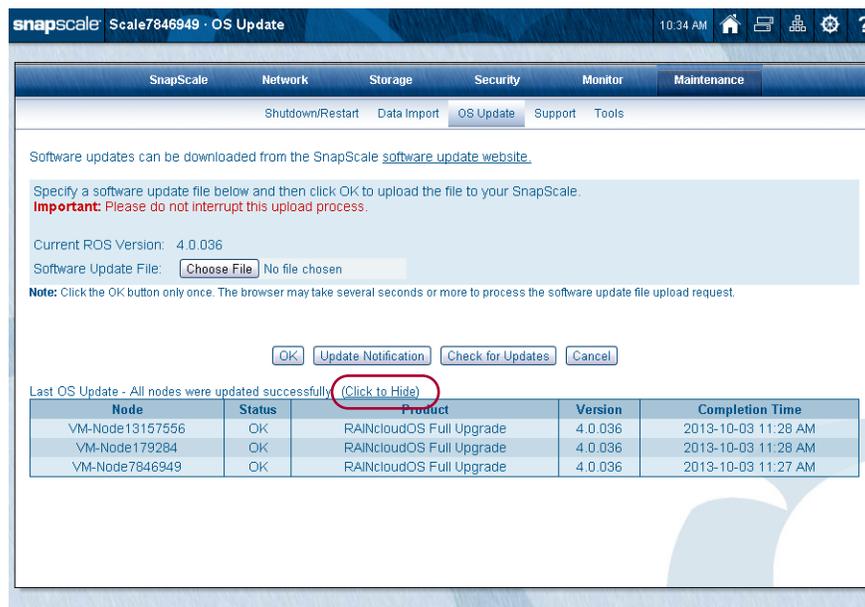


Information about the last RAINcloudOS update is listed at the bottom of the page, and may include the status of the update, product and version, and the completion time.



CAUTION: Do not interrupt the update process. You may severely damage the cluster's ability to run if you interrupt a software update operation.

To view the last OS update for each of the nodes, click the link under the buttons to show (or hide) the update table:



Update the RAINcloudOS Software

1. Click the **Check for Updates** button. If an update is available, follow the instructions on the page to download it.

NOTE: If the cluster does not have access to the Internet, download the latest RAINcloudOS image (.gsu) or other software package from the [Overland Storage website](#).

2. On the **OS Update** page, click **Browse**, locate the downloaded file, and select it.
3. Click **OK** to start the update (or **Cancel** to stop).

The software package is uploaded and then prompts you to reboot the cluster to perform the upgrade.

Note the following caveats:

- If upload fails on one or more nodes, the upgrade will abort with an error and list the nodes that had problems.
- After an upgrade and reboot, the **OS Update** page shows the status of the last update performed (success/failure) for each node.
- Snap EDR cannot be installed using the **OS Update** page.

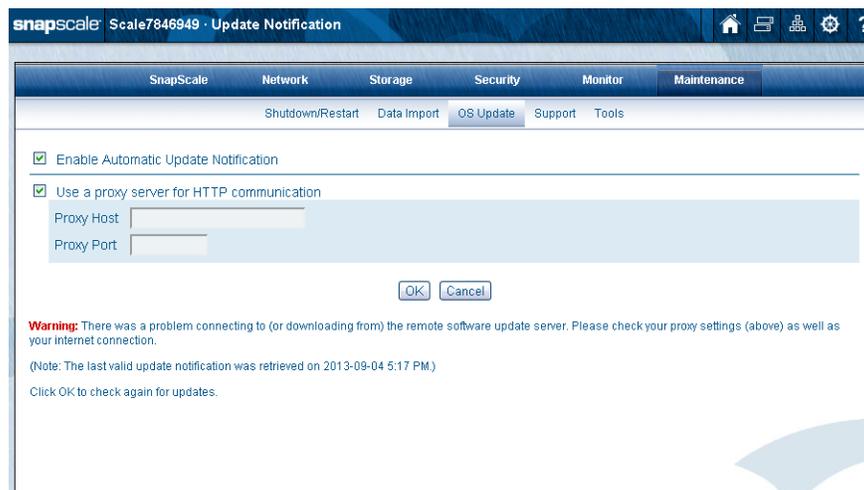
Software Update Notification

You can configure RAINcloudOS to display an alert when RAINcloudOS or Snap EDR updates are available for the cluster. When enabled, **Update Notification** checks weekly for RAINcloudOS or EDR updates that are applicable to the cluster. If updates are available, a banner alert is displayed just below the menu bar on all Web Management Interface pages.

NOTE: You can choose to hide the banner by clicking the [Remind me later](#) or [Hide this message](#) link on the banner. If [Remind me later](#), the Web Management Interface displays the banner after the next check for updates; if [Hide this message](#), the banner is hidden for the update in question until a later version is released.

Configuring Update Notification

1. Click the **Update Notification** button to display the options:



2. Check the **Enable Automatic Update Notification** option box.
3. If your environment requires using a proxy server for external web-based communication, check the **Use a proxy server for HTTP communication** check box and complete the **Proxy Host** and **Proxy Port** fields.
4. Click **OK**.

Manually Checking for Updates

Click the **Check for Updates** button to force the cluster to immediately search for applicable updates. If an update is available, it is displayed with information about it and a link to download the software.



Support

The **Support** page provides an easy way to contact Overland Technical Support.

 **IMPORTANT:** The **Support** page is not accessible until you have configured **Email Notification** in the **Tools** submenu.

Once email is configured, the **Support** page is available with your contact information entered:

Phone Home Support

Once email notification has been configured, Phone Home Support becomes available for registered SnapScale clusters. Phone Home Support optionally uploads and emails system logs and files that contain information useful for troubleshooting purposes to Overland Storage technical support. You can use the **Support** page under Maintenance to open a new case with technical support; or, in the course of working to resolve an issue, a technical support representative may ask you to fill out and submit this page. If a case is already in progress, you will need to enter the case number provided by the technical support representative.

NOTE: Phone Home Support interacts with two fields on the Email Notification page (Maintenance > Tools > Email Notification): (1) To use Phone Home Support, you must enter a valid SMTP server IP address on the Email Notification page; and (2) the first email address listed in the Recipients field populates the Reply-to Address field on the Support page.

Complete the following fields as appropriate, then click **OK**:

Text Field	Description
Subject	(Required) Enter a concise description that identifies the issue.
Case	(Required) Select New Case if you are emailing technical support for the first time. Select Existing Case if you have previously contacted technical support concerning the issue.
Case Number	If you selected Existing Case above, enter the case number provided by technical support.
Reply-to Address	(Required) This field defaults to the first email address entered as a recipient on the Email Notification page. Enter an email address that will serve as the contact address for this issue.
Cc Reply-to Address	To receive a copy of the email and system information attachment, check the Cc Reply-to Address box.
Comments	(Required) Enter additional information that will assist in the resolution of the problem.

NOTE: A Debug Logging option is located on this page but should only be changed at the direction of Overland Support. Click the link to change it.

Advanced Help

If you have an open case and have entered the **Case Number**, clicking the **Advanced** button opens additional options for the phone home feature. These options provide the ability to upload specific log files via FTP, which is sometimes necessary for the large logs the cluster can generate, and tech support may direct a user to use this for a particular case.

The screenshot shows the SnapScale Support interface. At the top, there's a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Below this is a sub-menu with Shutdown/Restart, Data Import, OS Update, Support, and Tools. The main content area contains a form for sending configuration details to Overland Storage Technical Support. The form includes fields for Subject (More Information), Case (Existing Case), Case Number (1234567), and Reply-to Address (user@overlandstorage.com). There's a checkbox for 'Cc Reply-to Address' and a large text area for Comments. Below the form is the 'Advanced Support Properties' section, which includes FTP settings (Server, Path, User, Password) and two radio button options for uploading information: 'Upload only information about nodes related to a specific file or directory' (selected) and 'Upload information about specific nodes'. A list of nodes is shown below, with 'VM-Node7846949 - 10.25.11.161 - (No description.)' selected. The dialog has 'OK' and 'Cancel' buttons at the bottom.

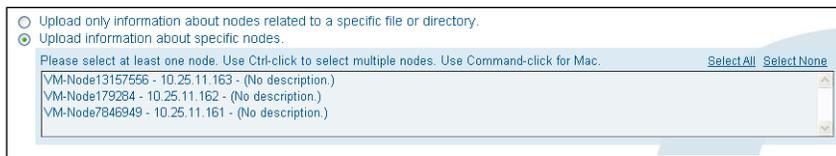
The additional options include:

- **FTP Server** – Name of the server to upload to.
- **FTP Path** – FTP server path used to upload to.
- **FTP User** – Name of the user to log in to the FTP server.
- **FTP Password** – FTP user's password.
- Click an **upload option**:
 - **Upload only information about nodes related to a specific file or directory**
Use the **Share** drop-down list to select the share and **File or Directory** path field to enter the path to a file or directory to gather logs. Only select **(Use absolute path.)** from the drop-down list when directed by Overland Support.

This close-up shows the selected radio button option: 'Upload only information about nodes related to a specific file or directory.' Below it, there's a text prompt: 'Please specify a share and path to a file or directory under the share. Select "Use absolute path" only as directed by Overland Support.' There are two input fields: 'Share' with a dropdown menu showing 'SHARE1' and 'File or directory' with an empty text box. The other radio button option, 'Upload information about specific nodes', is unselected.

- **Upload information about specific nodes**

Select one or more nodes to upload logs from them. Use the **Select All** and **Select None** options on the right as needed.

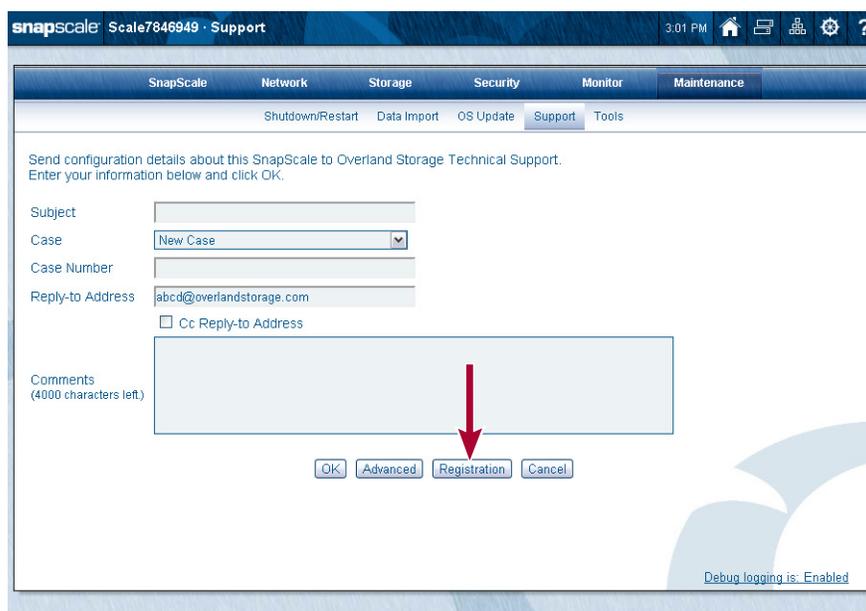


Registering Your Cluster

When you first start a cluster, a **Registration Reminder** page appears. Registering your cluster activates your warranty and allows you to create and track service requests. Registration also provides access to RAINcloudOS upgrades, third-party software, and exclusive promotional offers.

NOTE: Warranty information is available at <http://docs.overlandstorage.com/support>.

To register at a later time, use the **Registration** button on the **Support** page:



To Register Your Cluster

NOTE: To use this feature, access to the Internet is required.

Go to **Maintenance > Support** and click **Registration** to view the registration page:

Register your SnapScale to activate your warranty and stay informed of important software and product updates.
(Note: You must register your SnapScale whenever you add new nodes.)

Please enter your contact information below and then click the Download button to download the registration (text) file to your computer. You can then email the registration file to warranty@overlandstorage.com with the subject, **SnapScale Registration Request**

Name

Email Address

Company Name

Shipping Address (For RMA purposes.)

You can view the Overland Storage privacy policy by visiting www.snapserver.com/privacy (a new browser window will open).

Use this page to easily register your cluster:

1. Enter the **four required items** in the appropriate fields.
2. Click **Download Registration File**.
The information, including all the node data, is incorporated into a CSV file.
3. Depending on your browser settings, make sure you **save the file** to your local computer.
4. Email the downloaded registration file to warranty@overlandstorage.com.
Use the subject line **SnapScale Registration Request** for the email.

The same page is also used to update your registration information. For example, when you add new nodes to your cluster, they need to be added to the cluster registration so that they are also covered.

Once you have registered, you will receive a confirmation email to complete the registration.

Tools

The **Tools** option provides a submenu of general-purpose maintenance options and features.

Email Notification Configure the SnapScale to notify you when specific events occur.

Host File Editor Make your backup server known to this SnapScale.

Delete SnapScale Delete this SnapScale and all its data, and restore all nodes to their factory default state.

Email Notification

To configure the cluster to send email alerts in response to system events or activate Overland support, navigate to **Maintenance > Tools > Email Notification**.

The screenshot shows the 'Email Notification' configuration page in the SnapScale interface. The page title is 'Scale7846949 · Email Notification'. The interface includes a navigation bar with tabs for SnapScale, Network, Storage, Security, Monitor, and Maintenance. Under the Maintenance tab, there are sub-tabs for Shutdown/Restart, Data Import, OS Update, Support, and Tools. The main content area contains the following settings:

- Enable Email Notification
- SMTP Server: (Host name or IP address)
- SMTP Port: (Port number for SMTP server)
- Use Authenticated SMTP
- Use Secure Connection
- Email Address of Sender:
 - Use default: Scale7846949@devnet.myoverland.net
 - Use specific:
- Email Addresses of Recipients:
 - (optional)
 - (optional)
 - (optional)
- Send email notification for the following events:
 - Node shutdown/restart
 - Administrative operation event
 - Node hardware event
 - Storage usage warning event. (See SnapScale Properties)
 - SnapScale system event
- Send a test email to listed email addresses upon saving settings.

Buttons for 'OK' and 'Cancel' are located at the bottom of the form.

To set up email alerts, you need: (1) the SMTP server's IP address; and (2) the email address of each recipient who is to receive the alert.

Configuring Email Notification

Edit settings as described in the following table, and then click **OK**.

Option	Description
Enable Email Notification	To enable email notification, check the Enable Email Notification check box. Unchecking the box disables it.
SMTP Server	Enter a valid SMTP server IP address or host name.
SMTP Port	Enter a port number for the SMTP server or accept the default.
Use Authenticated SMTP	Check this box to authenticate when an email is sent to the SMTP server by the cluster. Provide an authentication User Name and Password in the fields that appear when the feature is enabled.
Use Secure Connection	Check this box to encrypt emails from the cluster. STARTTLS and TLS/SSL encryption protocols are supported.
Email Address of Sender:	Choose: <ul style="list-style-type: none"> The default address (<i>clustername@domain</i>) where <i>domain</i> is the DNS domain name (or the Management IP address (<i>clustername@ipaddress</i>) if no DNS domain name is configured. Specify a specific sender.
Recipients	Enter one or more email addresses to receive the notifications. One address is required. Three additional email addresses can be added.

Option	Description
Send Email Notification Events	Check the boxes next to the events you wish to be notified about: <ul style="list-style-type: none"> • Node shutdown/restart – A node shuts down or reboots due to an automatic or manual process. • Node hardware event – The internal temperature for a node exceeds its maximum operating temperature or other hardware problems. • SnapScale system (cluster) event – A change or error occurs that impacts the entire cluster. • Administrative operation event – A Data Import operation has finished or experienced an error. • Storage usage warning event – Storage space usage on a volume reaches either the maximum utilization or the critical utilization setting. See SnapScale Properties in Chapter 3.
Send a Test Email	To verify your settings, check Send a test email , then click OK.

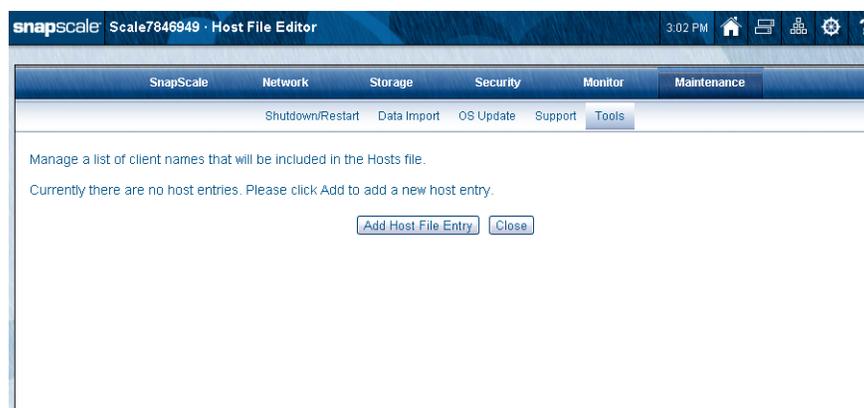
If **Send a Test Email** is checked, when you save your changes, the following email is sent to all configured email recipients:

```
Cluster Name: Scalennnnnnn
Node Name: Nodennnnnnn
IP Address:
Severity: Testing
Node Number: nnnnnnn
```

This is a test message.

Host File Editor

Use this page to identify external hosts in the SnapScale cluster's hosts file. This page allows you to supply a hostname-to-IP address mapping that persists across system reboots.



Click **Add Host File Entry**, complete the fields as described on the table below, and then click **Add Host File Entry** again.

Use this table to complete the options shown:

Option	Description
IP Address	The IP address of the external host.
Host Name	Enter the fully qualified hostname for the external host, using the format: <i>myserver.mydomain.com</i> . NOTE: Some applications may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.
Alias (optional)	Enter an optional abbreviated address for the external host, using the format: <i>myserver</i> . NOTE: Some applications may require that you enter either one or both of these fields. See the OEM documentation to determine requirements.

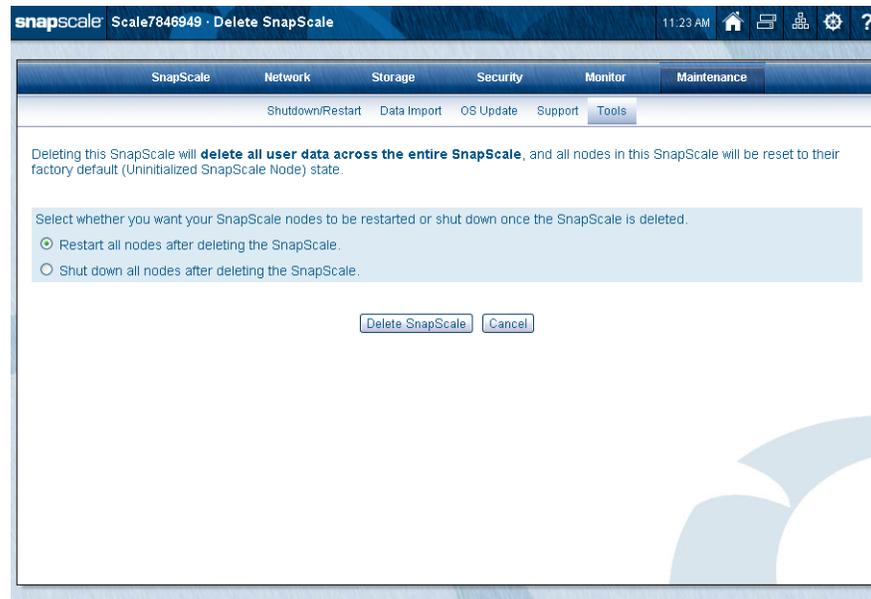
Delete SnapScale Cluster

This page is used to delete a SnapScale cluster and all its data.

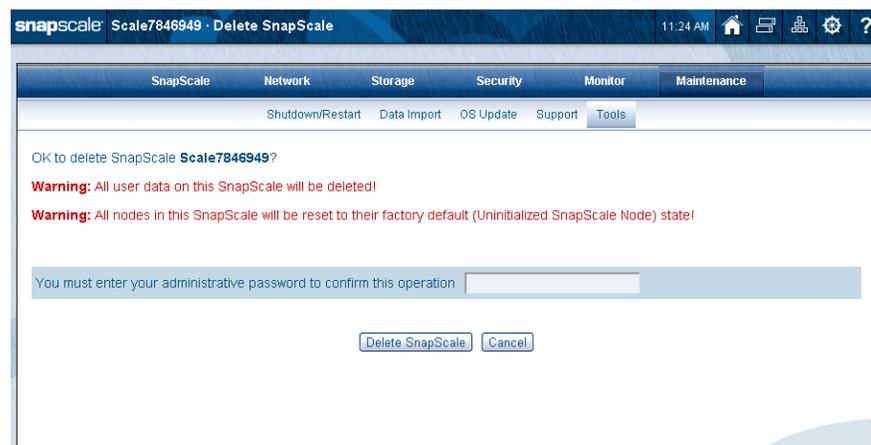


CAUTION: All data on all the nodes will be lost and all the nodes will be reset to their original factory default settings. No recovery is possible.

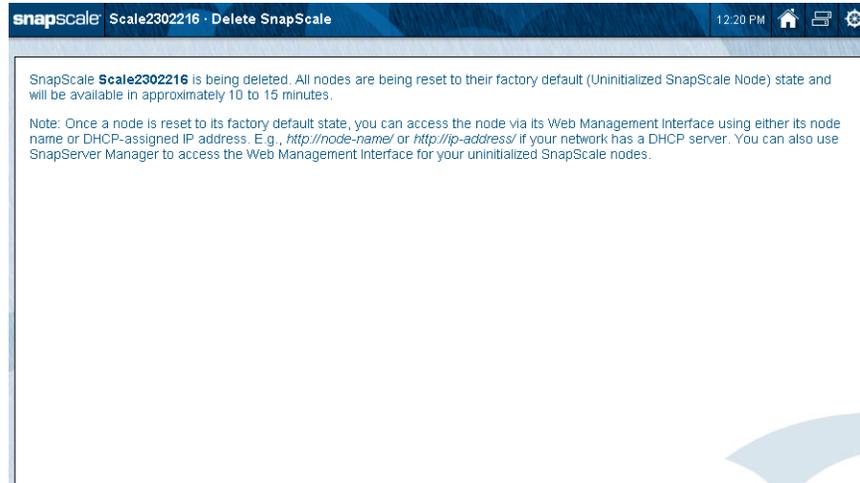
1. Select whether to restart or shut down all the nodes after the cluster is deleted.
 - **Restart all nodes** – After deleting the SnapScale cluster, the nodes reboot, they automatically perform a fresh install, and then they reboot as Uninitialized nodes.
 - **Shut down all nodes** – After deleting the SnapScale cluster, the nodes shut down. The next time the nodes are powered on, they automatically perform a fresh install and then reboot as Uninitialized nodes.



2. Click **Delete SnapScale**.
3. At the confirmation page, enter your **Admin password** and click **Delete SnapScale** again to start the process.

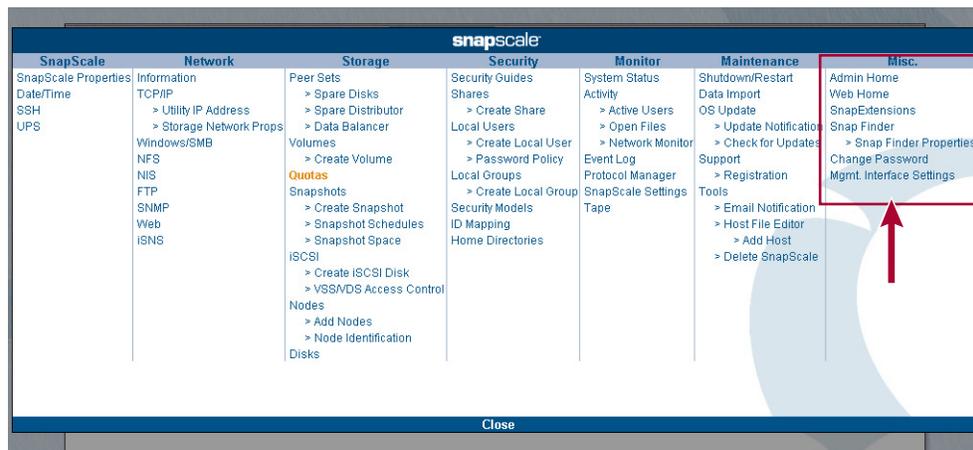


During the cluster deletion, an information page is shown (such as this one for the Restart option).



The RAINcloudOS site map (⚙️) provides links to all the web pages that make up the Web Management Interface. It also provides, in the last column, special links to higher level options and processes which is the focus of this chapter.

With the exception of **Mgmt. Interface Settings**, these options are also directly navigable from the various menus in the Web Management Interface. Also the **Home**, **Snap Finder**, **SnapExtensions**, **Site Map**, and **Help** options are accessible from any page by clicking their respective icon in the top right corner of the page (see the table in [Web Management Interface](#) in Chapter 2).



Topics in Misc. Options

- [Home Pages – Web/Admin](#)
- [SnapExtensions](#)
- [Snap Finder](#)
- [Change Password](#)
- [Management Interface Settings](#)

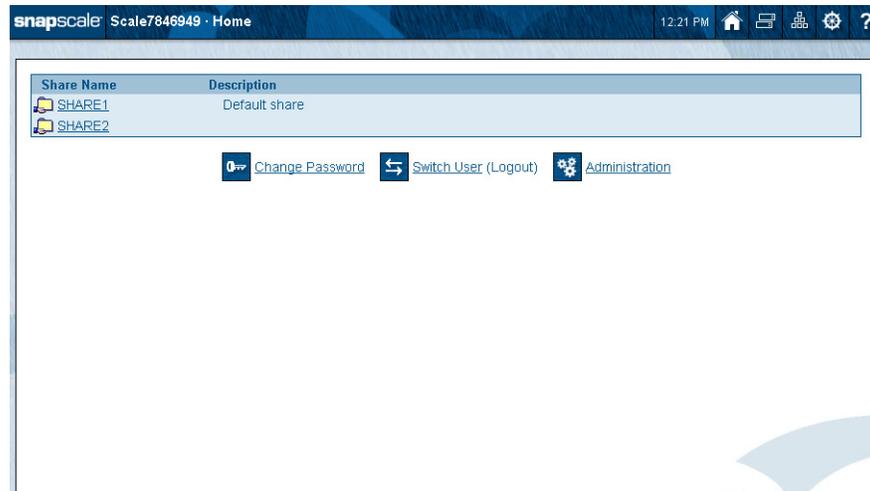
Home Pages – Web/Admin

When you first connect to the Web Management Interface (via http://cluster_mgt_ip/), the Web Home page (titled **Home**) is displayed. After clicking the **Administration** link (⚙️) and logging in, the Admin Home page (titled **Administration**) is displayed.

Once logged in, you can switch back and forth between the Web Home page and the Admin Home page using the **Home** (🏠) icon.

Web Home

The Web Home page shows a list of all the shares on the SnapScale cluster and has three links to key functions.



Clicking a share name/link displays the files and folders within the share. For users with admin rights, a key icon (🔑) appears next to each file/folder in the share. Clicking this icon displays a popup box with security information about the file/folder.

This page also provides three key administrative function links:

- **Change Password** (🔑) – Takes you to the **Change Password** page where you can change your administration password, or local users can change their password. Enter your **User Name** and **Current Password** for access. See [Change Password](#) on page 9-8.

Enter your user name, current and new passwords, and then click OK.

Important: Passwords are case-sensitive.

User Name

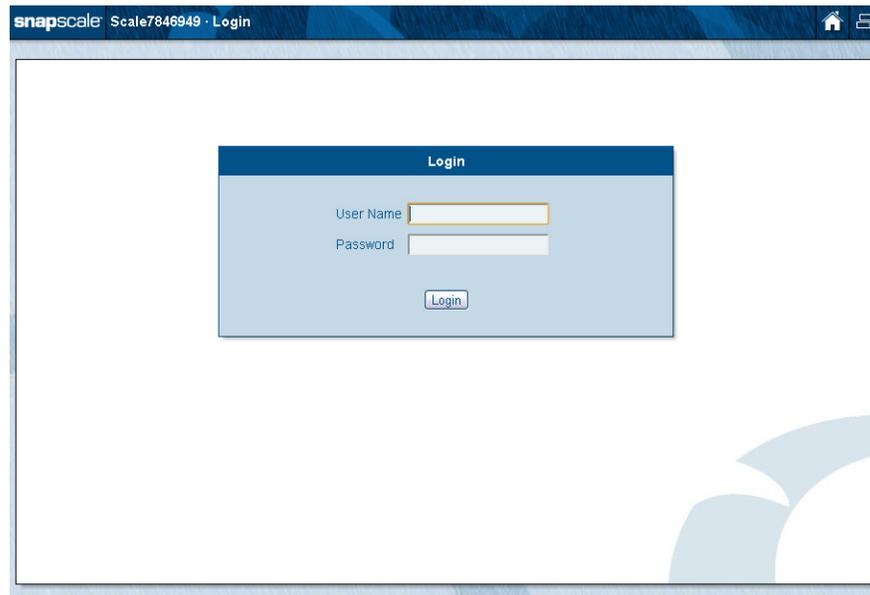
Current Password

New Password

Confirm New Password

OK Cancel

- **Switch User (Logout)** (↶) – Automatically logs out the current user and displays the **Login** page for the new user to gain access to the SnapScale cluster.



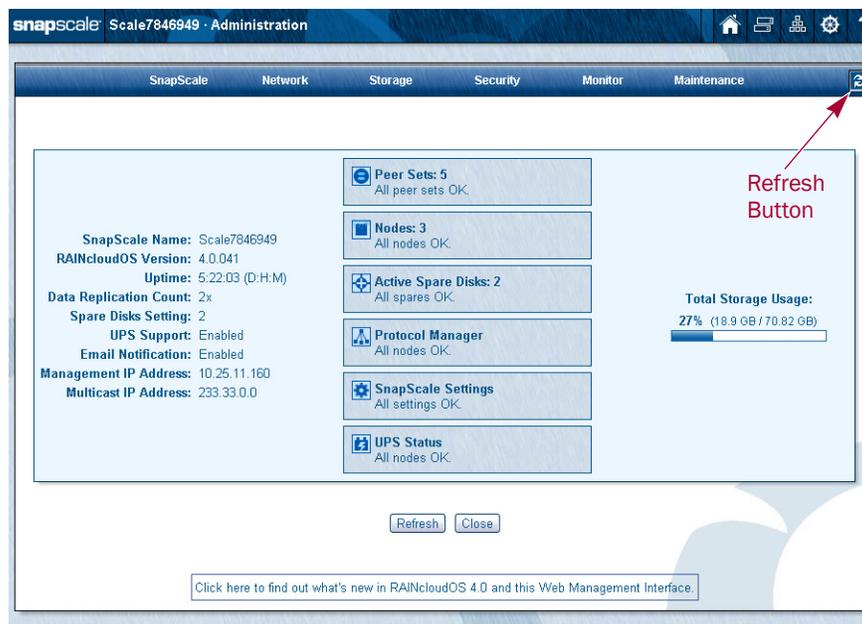
- **Administration** (⚙️) – Displays the **Administration** page (see [Administration on page 9-3](#)). You are prompted to log in if you have not already done so.

If any of the following conditions are present, you may not be able to access the Web Home page:

- **Require Web Authentication** is enabled (via **Network > Web > Require Web Authentication**) and you do not have a valid user name and password on the cluster.
- The cluster or node has not completed the **Initial Setup Wizard** (if this is the case, you will not be able to access the **Administration** page of the Web Management Interface either).
- **Web Root** is enabled (via **Network > Web > Enable Web Root**).

Administration

The **Administration** page is accessible by clicking the Admin Home link in the **Site Map**, or clicking either the **Administration** (⚙️) or **Home** (🏠) icons on the **Home** page, or by clicking the Home icon on any other page in the Web Management Interface. It provides a high-level view of the SnapScale status, links to key items such as peer sets and nodes, and a link to find out what's new in RAINcloudOS. The tabs at the top provide access to the various functions and features of the RAINcloudOS.



This table details the information on the Admin Home page:

Option	Description
SnapScale Name	Shows the name used to identify the cluster. The default name is "Scalennnnnn" (where nnnnnn is the appliance number of the node used to create the cluster).
RAINcloudOS Version	Lists the version of the RAINcloudOS running on all the nodes.
Uptime	Displays the length of time the cluster has been up and running since the last reboot.
Data Replication Count	Shows the data replication count establishing the level of data redundancy in the cluster (either 2x or 3x).
Spare Disks Setting	Displays the number of spare disks requested by user to automatically replace a failed peer set member.
UPS Support	Displays whether UPS support is enabled.
Email Notification	Displays whether Email Notification is currently enabled.
Management IP Address	Lists the IP address that is used to access the cluster through the Web Management Interface.
Multicast IP Address	Shows the multicast address used for inter-node messaging on the Storage network.
Peer Sets	Shows the total number of peer sets on the cluster and their status.
Nodes	Shows the number of nodes that make up the cluster and their status.
Active Spare Disks	Displays the total number of available hot spares and their status.
Protocol Manager	Provides the current status of networking protocols and the IP address assignment across the entire cluster.
SnapScale Settings	Provides the current status of any recently changed settings applied to the cluster.

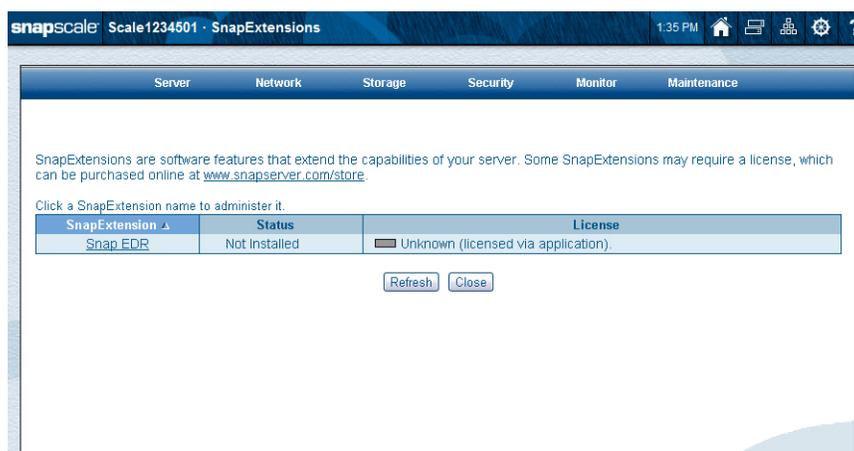
Option	Description
UPS Status	Displays UPS status for all nodes. This option is only displayed if UPS support is enabled.
Total Storage Usage	Displays a bar chart, percentage, and actual amount of cluster storage space used compared to total storage space available (used/total space). Mouseover also shows the free space.

The **Refresh** button at the bottom or the Refresh icon (🔄) on the right corner of the tab bar can be clicked to manually refresh the information.

From the **Administration** page, clicking 🏠 takes you to the Web Home page.

SnapExtensions

The SnapExtensions icon (🧩) opens the **SnapExtensions** page. This page is used to manage the SnapExtensions installed on your cluster.



If any SnapExtensions are installed, you can click the **SnapExtension** name in the table to display the management page for that extension.

Snap EDR

At the Snap EDR (Snap Enterprise Data Replicator) configuration page, you can install Snap EDR, configure it as either the Management Console or the Agent of another Management Console, or launch the Snap EDR Management Console interface. If configuring it as an Agent, enter the server name or cluster management name (<clustername>-mgt) of the Management Console.

Refer to [Configuring Snap EDR for RAINcloudOS](#) in [Appendix A](#) for details on installing and configuring Snap EDR.

NOTE: All other agents and management consoles must be able to resolve this cluster by the cluster management name (<clustername>-mgt) to the Management IP. This is typically done by a host record in the DNS.

After Snap EDR finishes its configuration, the **Management Console** page is shown.

For information on using Snap EDR, see [Snap Enterprise Data Replicator](#) in [Appendix A](#).

Snap Finder

Snap Finder () is a powerful tool that lists all the SnapScale clusters, Uninitialized nodes, and SnapServer appliances on your network (and on a remote network segment if so configured), and shows the current status of each. Click the unit name (if you have name resolution) or IP address of a cluster, node, or server to access it through the Web Management Interface.

NOTE: You can sort the columns (ascending or descending order) by clicking the column heading.



Server	Status	IP Address	OS Version	Model	Number	Avail Cap.	Total Cap.
MyCluster	Online	10.25.12.15	ROS 4.0.037	-	-	664.55 GB	1.06 TB
Node2300950	OK	10.25.2.143	ROS 4.0.037	X2	2300950	-	-
Node2300952	OK	10.25.2.29	ROS 4.0.037	X2	2300952	-	-
Node2300998	OK	10.25.2.1	ROS 4.0.037	X2	2300998	-	-
Node2303094	OK	10.25.2.127	ROS 4.0.037	X2	2303094	-	-
Node2303196	OK	10.25.2.132	ROS 4.0.037	X2	2303196	-	-
Node2303200	OK	10.25.2.195	ROS 4.0.037	X2	2303200	-	-
Node2413884	OK	10.25.6.19	ROS 4.0.037	X2	2413884	-	-
Node2413920	OK	10.25.2.193	ROS 4.0.037	X2	2413920	-	-
Node2414216	OK	10.25.2.248	ROS 4.0.037	X2	2414216	-	-
Node2414240	OK	10.25.2.237	ROS 4.0.037	X2	2414240	-	-
Scale2414274	Online	10.25.18.100	ROS 3.3.023	-	-	12.20 TB	12.21 TB
Scale2414338	Online	10.25.12.70	ROS 4.0.037	-	-	9.33 TB	9.33 TB
Scale3238585	Online	10.25.12.210	ROS 4.0.037	-	-	29.91 GB	30.81 GB
Scale7846949 (This SnapScale)	Online	10.25.11.160	ROS 4.0.036	-	-	51.99 GB	70.82 GB
sj-ae-620-test	OK	10.25.15.43	GOS 6.5.023	620	727636	159.47 GB	159.88 GB
sj-ae-620-test1	OK	10.25.15.31	GOS 6.5.023	620	724862	789.92 GB	1.09 TB
Eccles	Online	10.25.12.80	ROS 3.3.027	-	-	49.63 GB	50.85 GB
JH-VM1	OK	10.25.3.30	GOS 7.2.114	DX1	2198929	106.61 GB	107.01 GB
lambtron	OK	10.25.10.165	GOS 7.2.117	DX1	2301024	1.24 TB	2.00 TB
beryl	OK	192.168.48.166	GOS 7.2.124	DX1	2300028	4.20 TB	4.34 TB
bmcala1	OK	10.25.2.182	GOS 7.2.117	N2000	14391761	3.00 GB	6.09 GB
bmgs3	OK	10.25.3.19	GOS 7.2.123	N2000	10187012	6.08 GB	6.09 GB

The following table details the columns in the table:

Identification	Description
Server	Name of the SnapScale cluster, Uninitialized node, or SnapServer appliance. The default cluster name is <code>Scalennnnnn</code> , where <code>nnnnnn</code> is the number of the node originally used to create the cluster. For example, "Scale2302216."
Status	<ul style="list-style-type: none"> The status of the SnapServer or Uninitialized node (for example, OK or Fan Failure). The status of a SnapScale cluster is always Online.
IP Address	The IP address of the SnapServer, Uninitialized node, or the Management IP address of the SnapScale cluster.
OS Version	The OS version currently installed on the SnapServer, Uninitialized node, or SnapScale cluster.
Model	The hardware model number of the SnapServer or Uninitialized node. This field is not applicable to a SnapScale cluster.
Number	The server or node number derived from the MAC address of the primary Ethernet port, used as part of the default name. This field is not applicable to a SnapScale cluster.
Avail Cap	The available capacity on the cluster or SnapServer. This field is not applicable to an Uninitialized node.

Identification	Description
Total Cap	The total capacity on the cluster or SnapServer. This field is not applicable to an Uninitialized node.

To enable remote discovery of clusters, nodes, or servers on a different subnet or to display a warning icon for SnapServers or Uninitialized nodes with an enabled Ethernet port that has no link, click the **Properties** button to open the **Snap Finder Properties** page.

Edit Snap Finder Properties

Anyone with administrative privileges can view or edit the Snap Finder properties. Click the **Properties** button to access the server or node number page.

From this page you can select to display a warning icon for Uninitialized nodes or SnapServers with an enabled Ethernet port that has no link and enable remote discovery of units on a different subnet. Complete the following fields and then click **OK** to return to the **Snap Finder** page:

Option	Description
Display warning if any of a server's Ethernet ports have no link	Check to display a warning icon in the Status column for any nodes or SnapServers that have an enabled Ethernet port with no link. By default, this box is unchecked.
Enable Remote Server Discovery	Check to enable remote discovery of clusters, nodes, or SnapServers on a different subnet.
Add Server	Enter the host name or IP Address of a cluster, node, or server in the field to the right of the Add button, and click Add to incorporate it into the list of Remote Discovery Servers. Remote Discovery Servers send information about themselves as well as all other servers they've discovered on the remote network.
Delete Server	Select a cluster, node, or server, in the Remote Discovery Servers field and click Delete . When asked to confirm the deletion, click Delete again.

Change Password

To enhance the security of your SnapScale cluster, it is recommended that users change their passwords regularly. This is done using the **Change Password** page.



Enter your user name, current and new passwords, and then click OK.

Important: Passwords are case-sensitive.

User Name

Current Password

New Password

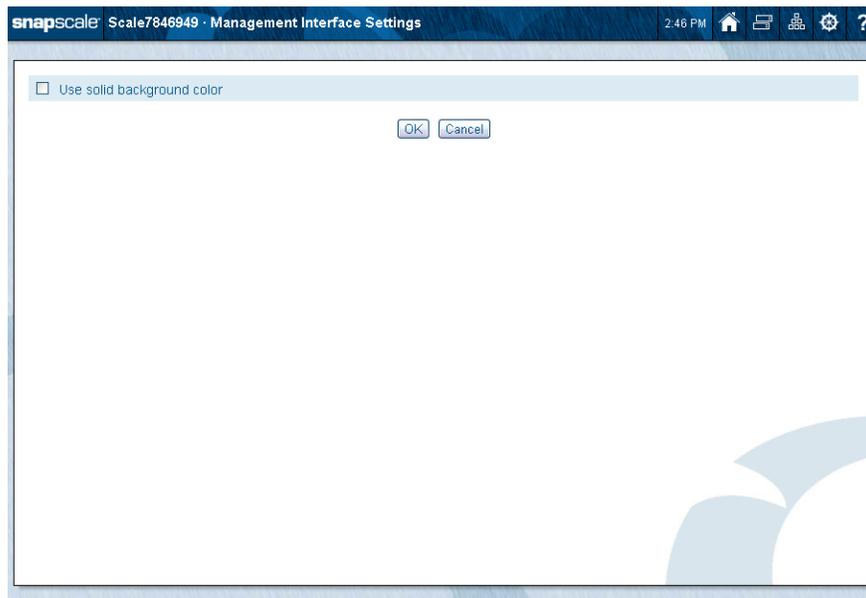
Confirm New Password

Changing Your Password

1. On the **Home** page, click the **Change Password** link ().
2. At the **Change Password** page, enter your **User Name** and **Current Password**.
3. Enter and confirm your **new password**.
Passwords are case-sensitive. Use up to 15 alphanumeric characters without spaces.
4. Click **OK**.
5. At the confirmation page, click **OK** again.
You are returned to the **Home** page.

Management Interface Settings

The Web Management Interface default background is light blue with the stylized “O” symbol. This can be changed to a solid blue background on the **Management Interface Settings** page by clicking the **Site Map** icon (⚙️), then choosing the **Management Interface Settings** link, and then checking the **Use solid background color** box. Click **OK** to make the change.



This appendix provides a brief description of the supported third-party backup solutions and the Snap Enterprise Data Replicator (EDR) software. Application Notes are available for installing and configuring third-party agents and media servers to work with SnapScale.

Topics in Backup Solutions:

- [Backup and Replication Solutions](#)
- [Snap Enterprise Data Replicator](#)
- [Backup via SMB or NFS](#)
- [Backup via Agent or Media Server](#)

Backup and Replication Solutions

RAINcloudOS supports several backup methods, including third-party off-the-shelf backup applications and applications that have been customized and integrated with RAINcloudOS on the SnapScale cluster:

- Data and security metadata backup and replication can be performed using the built-in Snap EDR.
- Backup over network file protocols can be performed using various backup packages that can access the cluster via SMB or NFS.
- Backup from the cluster or to a tape attached to a cluster node can be performed using supported backup agents and media libraries installed on one cluster node.

Snap Enterprise Data Replicator

Snap EDR provides server-to-server synchronization by moving, copying, or replicating the contents of a share from one cluster or server to another share on one or more different clusters or servers. It comes preinstalled on SnapScale clusters and activates a 45-day free trial if configured as a Management Console.

Snap EDR consists of a Management Console and a collection of Agents. The Management Console is installed on a central system. It coordinates and logs the following data transfer activities carried out by the distributed Agents:

- Replicates files between any two systems including SnapServers, SnapScale clusters, and Windows, Linux, and Mac Agents.
- Transfers files from one source host to one or more target hosts
- Transfers files from multiple hosts to a single target host, and stores the files on a local disk or locally attached storage device.
- Backs up data from remote hosts to a central host with locally-attached storage.

- Restores data from a central storage location to the remote hosts from which the data was originally retrieved.

Snap EDR Usage

The Snap EDR software distribution package comes preinstalled on the cluster and must first be installed in SnapExtensions and then configured before it's available for use. It operates under the cluster management name (in the style of `<clustername>-mgt`) and the Management IP address.

All other Snap EDR installations (including another machine running as the Management Console that the cluster registers to, other Agents that register to a Snap EDR Management Console running on the cluster, or other Agents replicating to/from the cluster) need to be able to resolve the cluster's management name to the cluster Management IP in order to interoperate properly with the cluster. This can be accomplished via a DNS host record or local hosts file entries.

When installing EDR to the cluster for the first time, it installs to all nodes and runs on whichever node currently serves as the Management node.

When adding a node to an existing cluster with Snap EDR installed, it is automatically installed on that node.

Configuring Snap EDR for RAINcloudOS

To configure the cluster as a Snap EDR Management Console or an Agent:

1. Click the **SnapExtensions** icon located in the upper right corner of the Web Management Interface.
2. If necessary, install the **software package**:
 - a. Run the **installation routine** from SnapExtensions.
SnapExtensions displays a Snap EDR link and the status **Not Installed**.
 - b. Click the **Snap EDR link** and confirm the installation.
Wait for the installation to complete. The **SnapExtensions** page then displays the **Snap EDR Configuration** link.
3. Click the **link** to launch the **Management Console/Agent** configuration page.
4. Select either the **Configure as the Management Console** or **Configure as the Agent** button.

NOTE: If you are configuring the cluster as an Agent, you must provide the server name (for a SnapServer) or cluster management name of the Management Console (for a SnapScale). The cluster must be able to resolve the name to the correct IP address.

5. Once the cluster is configured, select the following options from the page that appears:

Option	Description
Click here to configure jobs	Opens the Management Console where jobs can be scheduled.
Stop Service	Stops all services.

Option	Description
Restart Service	Restarts all services.

 **CAUTION:** Use only if you have encountered a problem, and customer support advises you to restart the service. Any jobs currently running will stop and will not resume when you restart the service.

Scheduling Jobs in Snap EDR

To schedule jobs, click the **Snap EDR** link in the **Site Map** (under **Misc.**).

For complete information on scheduling jobs in Snap EDR, see the *Snap EDR Administrator's Guide*.

Backup via SMB or NFS

A SnapScale cluster can be backed up via standard file server access.

In this configuration, the backup server is set up to use either SMB or NFS to connect to the cluster, examine the file system, and then back up the data onto itself. No special agents or media servers are needed.

Backup via Agent or Media Server

A SnapScale cluster can be backed up using a third-party Linux agent or media server installed on one cluster node. A tape drive or library can be attached to the node to provide local backup of the cluster.

Special Application Notes for installing the backup agent or media servers can be found on the Overland SnapScale Support website (<http://docs.overlandstorage.com/snapscale>).

NOTE: The backup packages shown in the Application Notes do not support the backup of Windows ACLs. If Windows ACL backup is critical, Overland Storage recommends using Snap EDR or creating a backup via SMB/CIFS to back up the cluster.

Utility IP Address

To provide continuous access to a specific cluster node with a backup agent or media server installed even when a public IP address changes, a special Utility IP address can be assigned that operates in addition to other cluster-assigned IP addresses and is always associated with the specific node. Then, if the public IP address of that node changes, backups continue to function without the need for the administrator to take action.

This IP address setting works for both the node as a backup source (backup from the node to the backup server) and the node as a backup target (backup from any machine including the node to an attached tape drive).

This appendix provides additional information and configuration options about accessing shares and files on the SnapScale cluster.

File and directory security can be configured using either Windows NTFS-style security or classic Unix-style security. The type of security present on a file or directory is its “security personality” (see [File-level Security](#) on [page B-4](#)). The default security model on newly-created volumes is always Windows/Mixed. It can be changed to a Unix security model if necessary.

RAINcloudOS on the SnapScale supports share-level as well as file- and directory-level permissions for all local and Windows domain users and groups (see [Windows ACLs](#) on [page B-4](#)).

Topics in Shares and File Access:

- [Security Model Rules](#)
- [Security Model Management](#)
- [Special Share Options](#)
- [File and Share Access](#)
- [File-level Security](#)

Security Model Rules

Files and directories are stored on the cluster on volumes with a configured “security model.” The security model on the volume governs the permitted security personalities, the default personalities, and the ability to change personalities on child files and directories.

Windows/Mixed Security Model:

- Files and directories created by SMB clients have the Windows security personality. Permissions will either be inherited according to the ACL of the parent directory (if Windows) or will receive a default ACL that grants the user full access only (if the parent is UNIX or has no inheritable permissions).
- Files and directories created by non-SMB clients will have the UNIX personality. UNIX permissions will be as set by the client (per the user’s local umask on the client).
- The security personality of a file or directory can be changed by any user with sufficient rights to change permissions or ownership. If a client of one security personality changes permissions or ownership of a file or directory of a different personality, the personality will change to match the personality of the client protocol (for example, if an NFS client changes UNIX permissions on a Windows file, the file will change to the UNIX personality).

UNIX Security Model:

- Files and directories created by non-SMB clients will have the UNIX personality. UNIX permissions will be as set by the client (per the user's local umask on the client).
- Files and directories created by SMB clients will have the UNIX personality. UNIX permissions are set to a default.
- The personality of files and directories cannot be changed on a UNIX security model. All files and directories always have the UNIX personality.

Security Model Management

A single security model can be set on a storage volume at the root level but different security models cannot be set on the directories immediately underneath the volume as the directories inherit the model from the top-level. Changes to a security model for the volume can optionally be propagated with the corresponding personality and default permission to all files and directories underneath.

When changing the security model:

- If changing from Windows/Mixed to UNIX, all files and directories will be changed to be owned by *admin* and *admingrp*, with UNIX permissions of *777(rwxrwxrwx)*.
- If changing from UNIX to Windows/Mixed, files and directories will be changed to default permissions that allow all users the ability to create and manage their own files and directories and to access other users' files and directories.

Special Share Options

The basic setup and configuration of shares on a SnapScale are handled on **Security > Shares** (see [Chapter 6, Security Options](#)). This section covers more details about the special options and features of share security on your SnapScale.

Hiding Shares

There are three ways a share can be hidden in RAINcloudOS:

- Name the share with a dollar-sign (\$) at the end. This is the traditional Windows method of hiding shares; however, it does not truly hide the share since Windows clients themselves filter the shares from share lists. Other protocols can still see dollar-sign shares.
- Hide the share from all protocols (except NFS) by:
 - While creating a share, navigating to **Security > Shares > Create Share > Advanced Share Properties** and checking the **Hide this Share** box.
 - Edit a share by selecting the share, clicking to expand **Advanced Share Properties**, and checking the **Hide this Share** box.

When a share is hidden this way, the share is invisible to clients, and must be explicitly specified to gain access.

NOTE: Hidden shares are not hidden from NFS, which cannot access invisible shares. To hide shares from NFS, consider disabling NFS access to the hidden shares.

- Disable individual protocol access to certain shares by:
 - While creating a share, navigating to **Security > Shares > Create Share > Advanced Share Properties** and enabling/disabling specific protocols.

- Edit a share by selecting a share, clicking to expand **Advanced Share Properties**, and enabling or disabling specific protocols.

Where to Place Shares

For security and backup purposes, it is recommended that administrators restrict access to shares at the root of a storage volume to administrators only. After initialization, all SnapScale clusters have a default share named *SHARE1* that points to the root of the default volume *Volume1*. The share to the root of the volume should only be used by administrators as a “door” into the rest of the directory structure so that, in the event that permissions on a child directory are inadvertently altered to disallow administrative access, access from the root share is not affected. This also allows one root share to be targeted when performing backups of the server. If it is necessary to have the root of the volume accessible, using the Hidden option helps ensure only those that need access to that share can access it.

File and Share Access

The shares feature controls access by users and groups. This section provides information on setting up the shares options to allow proper access to the files.

Cumulative Share Permissions

Share-level permissions on RAINcloudOS are applied cumulatively. For example, if the user “j_doe” has Read-Only share access and belongs to the group “sales”, which has Read/Write share access, the result is that the user “j_doe” will have Read/Write share access.

NOTE: Share-level permissions only apply to non-NFS protocols. NFS access is configured independently by navigating to [Security > Shares](#), selecting from the table the NFS Access level for the share, and modifying the client access as desired.

Snapshot Shares and On Demand File Recovery

A *snapshot share* is a read-only copy of a live share that provides users with direct access to versions of their files archived locally on the SnapScale via a snapshot. Users who wish to view or recover an earlier version of a file can retrieve it on demand without administrator intervention.

Snapshot shares are created during the course of creating a share. For instructions on accessing snapshot shares, see [“Configuring Share Access”](#) in [Chapter 6](#)

Creating a Snapshot Share

You create a snapshot share by going to **Security > Shares**. In the **Shares** column, click the share name. At the **Shares Properties** page, expand the **Advanced Share Properties** section, and check the **Create Snapshot Share** box.

For example, assume you create a share to a directory called *sales*, and you select the **Create Snapshot Share** option. When you connect to the server via a file browser or use the **Misc. > Web Home** link in the Site Map, two shares display:

```
SALES
SALES_SNAP
```

The first share provides access to the live volume, and the second share provides access to any archived snapshots. Other than read-write settings (snapshots are read-only), a snapshot share inherits access privileges from its associated live-volume share.

Accessing Snapshots Within the Snapshot Share

A snapshot share contains a series of directories. Each directory inside the snapshot share represents a different snapshot. The directory names reflect the date and time the snapshot was created.

For example, assume the snapshot share named *Sales_SNAP* contains the following four directories:

```
latest
2012-09-25.120000
2012-10-01.000100
2012-10-07.020200
```

The *latest* directory always points to the most recent snapshot (in this case, **2012-10-07.020200**, or October 7th, 2012, at 2:02 a.m.). A user may view an individual file as it existed at a previous point in time or even roll back to a previous version of the file by creating a file copy to the current live volume.

NOTE: The *latest* subdirectory is very useful for setting up backup jobs, as the name of the directory is always the same and always points to the latest available snapshot.

File-level Security

RAINcloudOS supports two “personalities” of filesystem security on files and directories:

- **Windows ACLs:** Windows NTFS-style filesystem permissions. Windows ACLs fully support the semantics of NTFS ACLs, including configuration, enforcement, and inheritance models (not including the behavior of some built-in Windows users and groups).
- **UNIX:** Traditional UNIX permissions (rwx) for owner, group owner, and other.

By default, volumes are created with the Windows/Mixed security model (Windows-style ACLs for files created by SMB clients and UNIX-style permissions for files created by other protocols and processes), and allow all users to create, delete, and configure permissions on their own files and to access files and directories created by other users.

Security Personalities and Security Models

The security personality of a file or directory is dependent on the security model of the root directory or volume in which the file or directory exists.

Files and directories in a Windows/Mixed security model can have either a Windows or UNIX security personality, depending on the network protocol used to create the file or change permissions on it. Files in a UNIX security model always have the UNIX security personality and can only be set by NFS clients.

Windows ACLs

RAINcloudOS fully supports Windows NTFS-style filesystem ACLs, including configuration, enforcement, and inheritance models. Inside Windows/Mixed root directories, files created and managed by Windows clients have the Windows security personality and behave just as they would on a Windows server. Clients can use the standard Windows 2000, 2003, XP, Vista, or Windows 7 interface to set directory and file permissions for local and Windows domain users and groups on the SnapScale.

Permissions are enforced for the specified users in the same manner for all client protocols, including non-SMB clients that normally have the UNIX security personality. However, if a non-SMB client changes permissions or ownership on a Windows personality file or directory (or deletes and recreates it), the personality will change to UNIX with the UNIX permissions specified by the client.

NOTE: Group membership of NFS clients is established by configuring the local client's user account or the NIS domain. Group membership of SnapScale local users or users ID-mapped to domain users is not observed by NFS clients. Therefore, ACL permissions applied to groups may not apply as expected to NFS clients.

Default File and Folder Permissions

When a file or directory is created by an SMB client, the owner of the file is the user who created the file (except for files created by local or domain administrators, in which case the owner is the **Administrators** group, mapped to the local **admingrp**), and the ACL will be inherited per the inheritance ACEs on the parent directory's ACL. The owner of a file or directory always implicitly has the ability to change permissions, regardless of the permissions established in the ACL. In addition, members of the SnapScale's local admin group, as well as members of Domain Admins (if the server is configured to belong to a domain) always implicitly have *take ownership* and *change ownership* permissions.

Setting File and Directory Access Permissions and Inheritance (Windows)

Access permissions for files and directories with the Windows security personality are set using standard Windows 2003, XP, Vista, 2008, or 7 security tools. RAINcloudOS supports:

- All standard generic and advanced access permissions that can be assigned by Windows clients.
- All levels of inheritance that can be assigned to an ACE in a directory ACL from a Windows client.
- Automatic inheritance from parent directories, as well as the ability to disable automatic inheritance from parents.
- Special assignment and inheritance of the CREATOR OWNER, CREATOR GROUP, Users, Authenticated Users, and Administrators built-in users and groups.

Procedure to set file and directory access permissions and inheritance in Windows:

1. Using a Windows 2003, XP, Vista, 2008, or 7 client, **map a drive** to the SnapScale, logging in as a user with change permissions for the target file or directory.
2. Right-click the file or directory, choose **Properties**, and then select the **Security** tab.
3. Use the **Windows security tools** to add or delete users and groups, to modify their permissions, and to set inheritance rules.

RAINcloudOS Ports

The following table lists the ports used by RAINcloudOS. The ROS Feature column lists access to the feature such as **Storage > iSCSI**. You can access the feature under the **Storage** tab in the **iSCSI** subsection of the Web Management Interface.

Port #	Layer	ROS Feature	Name	Comment
1	DDP		rtmp	Routing Table Management Protocol
1	TCP & UDP		tcpmux	TCP port service multiplexer
2	DDP		nbp	Name Binding Protocol
22	TCP & UDP	Server > SSH	ssh	Secure Shell (SSH) service
25	TCP & UDP	Server > Email Notification	smtp	Simple Mail Transfer Protocol (SMTP)
67	TCP & UDP	Network > TCP/IP	bootps	Bootstrap Protocol (BOOTP) services; also used by Dynamic Host Configuration Protocol (DHCP) services
68	TCP & UDP	Network > TCP/IP	bootpc	Bootstrap (BOOTP) client; also used by Dynamic Host Control Protocol (DHCP) clients
80	TCP & UDP	Web Management Interface	http	HyperText Transfer Protocol (HTTP) for World Wide Web (WWW) services
81	TCP	Web Management Interface	HTTP	Hypertext Transport Protocol
111	TCP & UDP	<ul style="list-style-type: none"> • Networking > NFS • Assist • SnapServer Manager 	sunrpc	Remote Procedure Call (RPC) Protocol for remote command execution, used by Network Filesystem (NFS) and SnapServer Manager
123	TCP & UDP	Server > Date/Time > Advanced	ntp	Network Time Protocol (NTP)
137	TCP & UDP	Network > Windows/SMB	netbios-ns	NETBIOS Name Services used in Red Hat Enterprise Linux by Samba
138	TCP & UDP	Network > Windows/SMB	netbios-dgm	NETBIOS Datagram Services used in Red Hat Enterprise Linux by Samba
139	TCP & UDP	Network > Windows/SMB	netbios-ssn	NETBIOS Session Services used in Red Hat Enterprise Linux by Samba
161	TCP & UDP	Network > Windows/SMB	snmp	Simple Network Management Protocol (SNMP)
162	TCP & UDP	Network > Windows/SMB	snmptrap	Traps for SNMP
389	TCP & UDP	Network > Windows/SMB	ldap	Lightweight Directory Access Protocol (LDAP)

Port #	Layer	ROS Feature	Name	Comment
443	TCP & UDP	<ul style="list-style-type: none"> • Web Management Interface • SnapServer Manager • SnapExtensions > Snap EDR 	https	Secure Hypertext Transfer Protocol (HTTP).
445	TCP & UDP	Network > Windows/SMB	microsoft-ds	Server Message Block (SMB) over TCP/IP
852	TCP	Network > NFS		Used by rpc.mountd
882	UDP	<ul style="list-style-type: none"> • Snap Finder • SnapServer Manager 	Sysbroker	Broadcast Discovery
933	UDP	Network > NFS		Used by rpc.statd
936	UDP	Network > NFS		Used by rpc.statd
939	TCP	Network > NFS		Used by rpc.statd
2005	TCP	SnapExtensions	SnapExtensions	Bridge from Servlet to Snap Extension framework
2049	TCP & UDP	Network > NFS	nfs [nfsd]	Network Filesystem (NFS)
2050	UDP	Network > NFS	mountd	
2051	UDP	Network > NFS	lockd	
2599	UDP	<ul style="list-style-type: none"> • Snap Finder • SnapServer Manager 	Sysbroker	Multicast Discovery
3052	TCP	Server > UPS		Port for monitoring UPS status
3205	TCP	Network > iSNS	iSNS	iSNS port
3260	TCP	Storage > iSCSI	iSCSI	iSCSI port
8001	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8002	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8003	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications
8005	TCP	Web Management Interface	tomcat	Tomcat Shutdown port
8008	TCP & UDP	Web Management Interface	http-alt	Tomcat - Apache Bridge
9049	TCP	Sysbroker		Sysbroker Shutdown Port
9050	TCP	Sysbroker		Sysbroker RPC Port
10001	TCP	Snap Extension	Snap Extension	Shutdown Port
32780	TCP	Web Management Interface	tomcat	Random Port
32781	TCP	Web Management Interface	tomcat	Random Port
49221	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications Port
49229	TCP	SnapExtensions > SnapEDR	SnapEDR	External Communications Port

Port #	Layer	ROS Feature	Name	Comment
1024 - 65535	TCP & UDP	Network > NFS Network > FTP	NFS FTP (Passive)	Dynamically allocated in runtime for user connections

Basic techniques for identifying and resolving common hardware and networking issues are described here.

Topics in Troubleshooting SnapScale Nodes

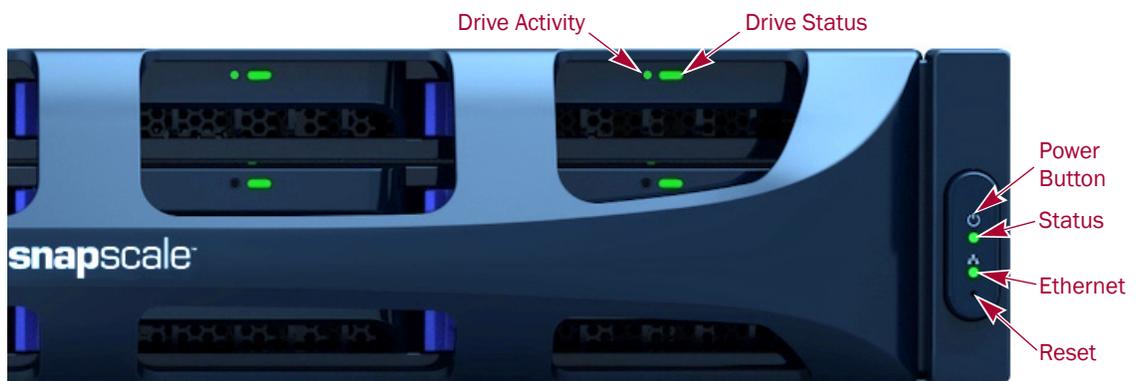
- [LED Indicators](#)
- [Network Reset](#)

LED Indicators

LED indicators provide information on the status of basic connectivity, disk drives, fan modules, and power supply modules.

SnapScale X2 Node LEDs

The SnapScale X2 has one network LED (Ethernet) and one system status LED on the Power Panel located on the right flange, along with a Power button and a Reset button. Each drive has two disk LEDs (Drive Activity and Drive Status) as shown in the following illustration:



The following tables describe the various states that may occur.

Drive Status LED

This is the oblong LED located on the center-right of each disk carrier. December 2013

Device State	LED State
No Disk Drive in Carrier	Off
Normal Operation	Solid green

Device State	LED State
Unit Identification Indicator	Flashing amber
Failed	Flashing red

Drive Activity LED

This is the round LED located on the center-left of each disk carrier.

Device State	LED State
Powered OFF / No Activity	Off
Drive Activity	Flashing green

Node Status LED

This is the round LED located just below the Power button on the right flange.

Device State	LED State
Powered OFF	Off
Booting	Solid amber
Normal Operation	Solid green
Shutting down	Flashing green
Maintenance Mode	Flashing green/amber

Client Network LED

This is the round LED located just below the Status LED on the right flange.

Device State	LED State
Powered OFF	Off
Link Up (SnapScale Powered ON)	Solid green
Link Down	Off

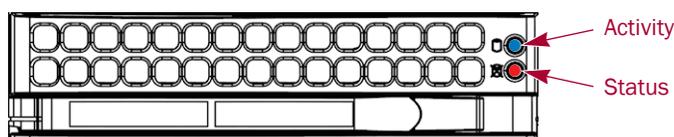
Power Supply Status LED

This is the LED located just above and to the right of the socket.

Device State	LED State
Normal Operation	Solid green
Standby	Solid red
Power Failure	Off
Fan Failure	Blinking red

SnapScale X4 Node LEDs

Each drive has two disk LEDs (top is Drive Activity and bottom is Drive Status):



LED	Description
Activity	Solid blue when a drive is present in the carrier. Blinks when there is drive activity.
Status	Red LED that: <ul style="list-style-type: none"> • Solid red when a drive fails. • Blinks red with all the other drive Status LEDs when node ID is activated. • Off if drive is OK or not present.

The Control Panel on the left flange has a series of LEDs as shown here:

Control Panel	Icon & Name	Description
	Power Button	This is the main power button. Turning off system power with this button removes the main power but keeps standby power supplied to the system. The power cords should be unplugged before service.
	Status	Used as follows: <ul style="list-style-type: none"> • Off when the node is off. • Double flashing green when booting. • Solid green during normal operation. • Flashing green when powering down. • Alternating blink/double blink green in maintenance mode.
	HDD	Always off.
	NIC 1	Indicates network link on the Ethernet 1 port when green .
	NIC 2	Indicates network link on the Ethernet 2 port when green .
	Overheat/ Fan Fail	When this LED flashes, it indicates a fan failure. When on continuously, it indicates an overheat condition.
	Power Supply	Used as follows: <ul style="list-style-type: none"> • Solid green during normal operation. • Solid amber or off if a module has failed, is not connected, or the node has been turned off.

Network Reset

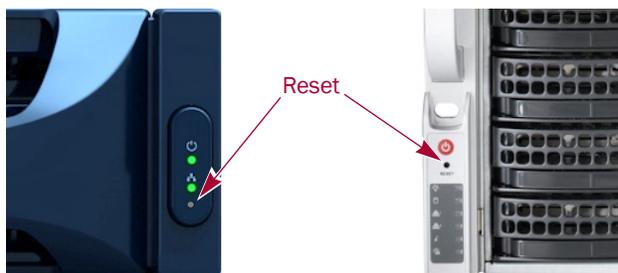
NOTE: The reset button is only operational for Uninitialized SnapScale nodes (not part of a cluster).

If an Uninitialized SnapScale node has been configured with incorrect network settings, the settings can be reset to the default values via the reset button.

The **Reset** button is accessed via a small hole on the side flanges:

- On the **X2** node, it is just below the Network LED on the Power panel located on the right flange.
- On the **X4**, the hole is located just below the Power button on the Power panel located on the left flange.

Verify that the server is fully booted (as indicated by the system/status LED) and, using the end of a straightened paper clip or the fine point of an instrument, press the **Reset** button to initiate a network reset.



The system will reboot after about a minute and the network settings will be reset to DHCP mode.

Master Glossary & Acronym List

NOTE: This is a general Overland Storage glossary and acronym list. Not all items may be found in this document or be used by this product.

1000BASE-T

1000BASE-T (also known as IEEE 802.3ab) is a standard for gigabit Ethernet over copper wiring. It requires, at a minimum, Category 5 cable (the same as 100BASE-TX), but Category 5e (Category 5 enhanced) and Category 6 cable may also be used and are often recommended. 1000BASE-T requires all four pairs to be present and is far less tolerant of poorly installed wiring than 100BASE-TX.

Access Permissions

A rule associated with a share, a file, or a directory on a disk drive to regulate which users can have access to the share and in what manner.

Address

An address is a data structure or logical convention used to identify a unique entity, such as a particular process or network device.

ACL

Short for *Access Control List*. A mechanism for restricting access to disk drive directories and files. ACLs are lists of access rights assigned to users and groups on files and directories.

ADS

Short for *Active Directory Service*. The preferred authentication method for Windows network users. This allows Active Directory users to connect to shares on the SnapServer or SnapScale cluster. The SnapServer and SnapScale both support the Microsoft Windows 2000 family of servers that run in native ADS mode.

Agent

A program that performs some information-gathering or processing task in the background. Remote agents provide the ability of a third-party server to communicate and control storage devices.

Algorithm

A sequence of steps designed to solve a problem or execute a process.

AllLocalUsers Group

The default group for all local users on SnapServers and SnapScale clusters. Local users are set up by the administrator. Network users or Windows domain users are not part of the AllLocalUsers group.

AllUsers Group

A collection of all users. The SnapServer or SnapScale cluster automatically maintains the AllUsers group.

ATA

Short for *Advanced Technology Attachment*. A standard interface for connecting storage devices to a PC.

Authentication

The validation of a user's identity by requiring the user to provide a registered login name and corresponding password.

Autonegotiation

An Ethernet feature that automatically negotiates the fastest Ethernet speed and duplex setting between a port and a hub or switch. This is the default setting and is recommended.

Autosensing

An Ethernet feature that automatically senses the current Ethernet speed setting.

Bar Code

The machine-readable representation of a product code. Bar codes are read by a scanner that passes over the code and registers the product code. The width of black lines and white spaces between varies. Combinations of lines and spaces represent characters. Overland uses 3-of-9 code (Code 39) where each character is represented by 9 bars, 3 of which are wide.

Bonding

A technology that treats two or more ports as a single channel, with the network using one IP address for the server. SnapServers and SnapScale clusters support load balancing and failover bonding modes.

Bus or Channel

A common physical path composed of wires or other media, across which signals are sent from one part of a computer to another. A channel is a means of transferring data between modules and adapters, or between an adapter and SCSI devices. A channel topology network consists of a single cable trunk that connects one workstation to the next in a daisy-chain configuration. All nodes share the same medium, and only one node can broadcast messages at a time.

CA

Short for *Certificate Authority*. A trusted third-party in a network that issues and manages security credentials.

Cat 5 Cable

Short for *Category 5*, it is network cabling that consists of four twisted pairs of copper wire terminated by 8P8C modular connectors. CAT 5 cabling supports frequencies up to 100 MHz and speeds up to 100 Mbps. (CAT 5e cabling supports frequencies up to 1000 MHz and speeds up to 1000 Mbps.) It can be used for ATM, token ring, 1000BASE-T, 100BASE-T, and 10BASE-T networking.

Cat 5 is based on the EIA/TIA 568 Commercial Building Telecommunications Wiring Standard developed by the Electronics Industries Association as requested by the Computer Communications Industry Association in 1985.

Cat 6 Cable

Short for *Category 6*, it is network cabling that consists of four twisted pairs of copper wire terminated by 8P8C modular connectors made to higher standards that help reduce noise caused by crosstalk and system noise. The ANSI/TIA-568-B.2-1 specification states the cable may be made with 22 to 24 AWG gauge wire, so long as the cable meets the specified testing standards.

It is designed for Gigabit Ethernet that is backward compatible with the Category 5/5e and Category 3 cable standards. Cat 6 features more stringent specifications for crosstalk and system noise. The cable standard provides performance of up to 250 MHz and is suitable for 10BASE-T, 100BASE-TX, and 1000BASE-T (Gigabit Ethernet).

Channel

A communications path between two computers or devices.

Checksum

The result of adding a group of data items that are used for checking the group. The data items can be either numerals or other character strings treated as numerals during the checksum calculation. The checksum value verifies that communication between two devices is successful.

CIFS

Short for *Common Internet Filesystem*. Also known as **SMB**. The default Windows protocol for communication between computers. A specification for an Internet file access protocol that complements HTTP and FTP.

daemon

A process that runs in the background.

default gateway

The router used when there is otherwise no known route to a given subnet.

degraded

A RAID state caused by the failure or removal of a disk drive in which data is consistent, but there is no redundancy.

DHCP

Short for *Dynamic Host Configuration Protocol*. A communications protocol that lets network administrators centrally manage and automate the assignment of IP addresses on a computer network. Each system that connects to the Internet/intranet needs a unique IP address.

Disaster Recovery

A strategy that allows a company to return to normal activities after a catastrophic interruption. Through failover to a parallel system or by restoration of the failed system, disaster recovery restores the system to its normal operating mode.

DNS

Short for *Domain Name Service*. A network service that translates domain names into IP addresses using a server that maintains a mapping of all host names and IP addresses. Normally, this mapping is maintained by the system administrator, but some servers support dynamic mappings.

Domain

A set of network resources in Windows 2000/2003/2008, such as users and groups of users. A domain may also include multiple servers on the network. To gain access to these network resources, the user logs into the domain.

Domain Name

The ASCII name that identifies the domain for a group of computers within a network.

Ethernet

The most widely installed LAN technology. 100BASE-T Ethernet provides transmission speeds of up to 100 Mbps. Fast Ethernet or 1000BASE-T provides transmission speeds up to 1000 Mbps and is typically used for LAN backbone systems, supporting workstations with 100BASE-T cards. Gigabit Ethernet (GbE) provides an even higher level of backbone support at 1000 Mbps (one Gigabit or one billion bits per second).

Ethernet Address

The unique six-digit hexadecimal (0-9, A-F) number that identifies the Ethernet interface.

Ethernet Port

The port on a network card to provide Ethernet access to the computer.

Event

Any significant occurrence or error in the system that may require notifying a system administrator or adding an entry to a log.

Expansion Slot

Area in a computer that accepts additional input/output boards to increase the capability of the computer.

Failover

A strategy that enables one Ethernet port to assume the role of another port if the first port fails. If a port fails on a server, the second port assumes its network identity (if the two Ethernet cards have been configured for failover). When the port comes back online, the original identities are restored. Failover is possible only in a multi-Ethernet configuration.

Failover/Failback

A combination of Failover and Failback. When a preferred path becomes unavailable, another path is used to route I/O until the preferred path is restored. In this case I/O will “fail back” to the preferred path once it is available again.

Fibre Channel

Fibre Channel (FC) is a gigabit-speed network technology which transports SCSI commands over Fibre Channel networks. Fibre Channel was primarily concerned with simplifying the connections and increasing distances, but later designers added the goals of connecting SCSI disk storage, providing higher speeds and far greater numbers of connected devices.

Filesystem

A type of data store which can be used to store, retrieve, and update a set of files.

Firmware

Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first switched on.

FTP

Short for *File Transfer Protocol*. A standard Internet protocol that provides a way to exchange files between computers on the Internet. By default, a SnapServer or SnapScale cluster is set up to be an FTP server.

Full-duplex

A type of transmission that allows communicating systems to both transmit and receive data simultaneously.

Gateway

The hardware or software that bridges the gap between two network subnets. It allows data to be transferred among computers that are on different subnets.

Gigabit Ethernet

Also known as GigE or GbE, this Ethernet standard uses a one Gigahertz (1000 Hz) clock rate to move data.

GID

Short for *Group Identification*. On a SnapServer, the unique ID assigned to each group of users for security purposes.

GSU

An image file used to upgrade the OS.

HBA

Short for *Host Bus Adapter*. An HBA is an I/O adapter that sits between the host computer's bus and the Fibre Channel loop and manages the transfer of information between the two channels. In order to minimize the impact on host processor performance, the HBA performs many low-level interface functions automatically or with minimal processor involvement.

Half-duplex

A type of transmission that transfers data in one way at a time.

Hidden Share

A share that restricts the display of the share via the Windows (SMB), Web View (HTTP/HTTPS), FTP, and AFP protocols. See also [SMB](#).

Host Name

The unique name by which a computer is known on a network. It is used to identify the computer in electronic information interchange.

Hot Swapping

The ability to remove and add disk drives to a system without the need to power down or interrupt client access to filesystems. Not all components are hot-swappable. Please read installation and maintenance instructions carefully.

HTTP

Short for *Hypertext Transfer Protocol*. An application protocol for transferring files (text, graphic images, sound, video, and other multimedia files) over TCP/IP on the World Wide Web.

HTTPS

Short for *Hypertext Transfer Protocol Secure*. The HTTP protocol using a Secure Sockets Layer (SSL). SSL provides data encryption, server authentication, message integrity, and client authentication for any TCP/IP connection.

Inheritance

In Windows permissions, inheritance is the concept that when permissions for a folder are defined, any subfolders within the defined folder inherit its permissions. This means an administrator need not assign permissions for subfolders as long as identical permissions are desired. Inheritance greatly reduces administrative overhead and also results in greater consistency in access permission management.

Initiator Device

An iSCSI system component that originates an I/O command over an I/O bus or network. An initiator issues the commands; a *target* receives them.

An initiator normally runs on a host computer. It may be either a software driver or a hardware plug-in card, often called a Host Bus Adapter (HBA). A software initiator uses one of the computer's Ethernet ports for its physical connection, whereas the HBA will have its own dedicated port.

Software initiators are readily available for most host operating systems. Hardware initiators are not widely used, although they may be useful in very high performance applications or if 10 Gigabit Ethernet support is required.

I/O (Input/Output)

The operation of transferring data to or from a device, typically through an interface protocol like CIFS, NFS, or HTTP. The SnapServer presents a filesystem to the user and handles block I/O internally to a RAID array.

IP

Short for *Internet Protocol*. The unique 32-bit value that identifies the location of the server. This address consists of a network address, optional subnetwork address, and host address. It displays as four addresses ranging from 1 to 255 separated by periods.

IQN

Short for *iSCSI Qualified Name*. A name format used in the iSCSI protocol. Initiators and targets have IP addresses, just like any other network entity. They are also identified using an iSCSI name, called the iSCSI Qualified Name (IQN). The IQN should be unique worldwide. It is made up of a number of components, specifying the date, identifying the vendor in reverse format, and then uniquely identifying the initiator or target. An example of an IQN is:

```
iqn.2001-04.com.example:storage:diskarray-sn-123456789
```

Since these IQNs are rather unwieldy, initiators and targets also use short, user friendly names (sometimes called alias names or just aliases).

iSCSI

Short for *Internet SCSI*. iSCSI is an IP-based storage networking standard for linking data storage facilities. iSCSI is a standard that defines the encapsulation of SCSI packets in TCP and then routing it using IP. It allows block-level storage data to be transported over widely used IP networks.

iSNS Server

Short for *Internet Storage Name Service Server*. A protocol enabling the automatic discovery, configuration, and management of iSCSI devices on a TCP/IP network.

Kerberos

A secure method for authenticating a request for a service used by ADS. Kerberos lets a user request an encrypted “ticket” from an authentication process that can then be used to request a service from a server. The user credentials are always encrypted before they are transmitted over the network.

In Windows 2000/XP, the domain controller is the Kerberos server. The Kerberos key distribution center (KDC) and the origin of group policies are applied to the domain.

LACP

Link Aggregation Control Protocol provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).

LAN

Short for *Local Area Network*. A network connecting computers in a relatively small area such as a building.

LCD

Short for *Liquid Crystal Display*. An electronic device that uses liquid crystal to display messages.

LED

Short for *Light-Emitting Diode*. An LED is a type of diode that emits light when current passes through it. Visible LEDs are used as indicator lights on electronic devices.

Linux

A UNIX-like operating system that was designed to provide personal computer users a free or very low-cost operating system comparable to traditional and usually more expensive UNIX systems.

Load Balancing

A process available only in multi-Ethernet configurations. The Ethernet port transmission load is distributed among two or more network ports (assuming the cards are configured for load balancing). An intelligent software adaptive agent repeatedly analyzes the traffic flow from the server and distributes the packets based on destination addresses.

Local Group/Local User

A group/user defined locally on a SnapServer or SnapScale cluster using the Web Management Interface. The local user is defined by the server or cluster administrator. Windows domain, ADS, and NIS users are not considered local.

MAC Address

Short for *Media Access Control address*, a hardware address that uniquely identifies each node of a network. In the Open Systems Interconnection (OSI) model, one of two sublayers of the Data Link Control layer concerned with sharing the physical connection to the network among several computers. Each Ethernet port has a unique MAC address. SnapServers with dual-Ethernet ports can respond to a request with either port and have two unique MAC addresses.

Maintenance Mode

A series of HTML pages in the Web Management Interface that allows you to perform repair, upgrade, or reinstall the OS in a disaster recovery situation.

MD5 Algorithm

MD5 is a way to verify data integrity, and is much more reliable than checksum and many other commonly used methods.

MIB

Short for *Management Information Base*. A formal description of a set of network objects that can be managed using the Simple Network Management Protocol (SNMP). The format of the MIB is defined as part of SNMP.

Mirroring

Used in RAID 1 and 10, a process of storing data on one disk and copying it to one or more disks, creating a redundant storage solution. RAID 1 is the most secure method of storing mission-critical data.

Mounted

A filesystem that is available.

MPIO

Short for *Multipath Input/Output*. A multipath solution built into Microsoft server-grade iSCSI operating systems.

MTU

Short for *Maximum Transfer Unit*. It is the largest size packet or frame, specified in octets (eight-bit bytes), that can be sent in a packet- or frame-based network.

Multihomed

A SnapServer that is connected to two or more networks or has two or more network addresses.

NFS

Short for *Network Filesystem*. A client/server application that allows a computer user to view and optionally store and update files on a remote computer as though they were on the user's own computer. The user's system needs to have an NFS client and the other computer needs the NFS server. The SnapServer and SnapScale are configured as an NFS server by default.

NIC

Short for *Network Interface Card*. A board that provides network communication capabilities to and from a computer.

NIS

Short for *Network Information Service*. A network naming and administration system for smaller networks that was developed by Sun Microsystems. NIS+ is a later version that provides additional security and other facilities. The SnapServer and SnapScale accept NIS users and groups.

NTFS

Short for *New Technology File System*. The standard file system used by Windows NT and later versions of the Windows operating system.

NTP

Short for *Network Time Protocol*. A protocol for synchronizing the system clocks of computers over a packet-switched network.

NVRAM

Abbreviation of *Non-Volatile Random Access Memory*, a type of memory that retains its contents when power is turned off.

Orphan

A disk drive that has become disconnected from its RAID either by accidental removal of the drive or the intermittent failure of the drive.

Parity

Error correction data. RAID5 and RAID6 store equal portions of each file on each disk and distributes parity information for each file across all disks in the group. This distributed parity allows the system to recover from a single disk drive failure.

Permissions

A security category, such as no access, read-only, or read-write, that determines what operations a user or group can perform on folders or files.

PoP

Short for *Proof of Purchase*. The number used to obtain a license key for an upgrade to third-party applications.

Portal

A target's IP address together with its TCP port number used in iSCSI systems.

POSIX

Short for *Portable Operating System Interface*. A set of standard operating system interfaces based on the UNIX operating system. The need for standardization arose because enterprises using computers wanted to develop programs that could run on multiple platforms without the need to recode.

Protocol

A standardized set of rules that specifies the format, timing, sequencing, and/or error checking for data transmissions.

PTP

Short for *Point-to-Point*. PTP is the common mode of attachment to a single host. PTP is sometimes used to attach to a Fibre Channel switch for [SAN](#) connectivity.

Quota

A limit on the amount of storage space on a volume that a specific user or NIS group can consume.

RAID

Short for *Redundant Array of Independent Disks*. A data storage scheme where multiple hard drives are combined to form a single logical unit which is highly reliable and gives good performance. Reliability is achieved by mirroring (the copying of data to more than one disk), striping (the splitting of data across more than one disk) and error correction (redundant data is stored to enable faults to be detected and corrected).

RAID 0 (Striped)

RAID 0 is ideal for environments in which performance (read and write) is more important than fault tolerance, or you need the maximum amount of available drive capacity in one volume.

Data is striped across multiple disks so that it can be read and written in parallel. It provides higher performance than a single disk, especially when reading or writing large files, but it is vulnerable to a disk failure. If any disk in the pool fails, the entire pool is effectively lost. For this reason, RAID 0 pools should only be used in cases where the loss of the data is unimportant, for example, because it can easily be recreated from another data source. The capacity of a RAID 0 pool is equal to the total capacity of all the disks making up the pool¹. For example, a RAID 0 pool made up of 4 x 1 TB disks will have a capacity of 4 TB.

RAID 1 (Mirrored)

RAID 1 is useful for building a fault-tolerant system or data volume, providing excellent availability without sacrificing performance. However, you lose 50 percent of the assigned disk capacity.

RAID 1 is also called disk mirroring: data is stored on two identical disks, so that if one disk fails, the other can still be used to access the data. Write operations are performed in parallel to both disks, so write performance is identical to that of a single disk; read operations can be done to either disk, so effectively read performance is doubled.

If one of the disks fails, it should be replaced. When it is replaced, the RAID pool will automatically be rebuilt by copying all the data from the surviving disk to the new disk. While the rebuild is occurring, there will be a degradation in performance.

Because disks are mirrored, the usable capacity of a pair of RAID 1 disks is only equal to the capacity of a single disk, so that a RAID 1 pool made of 2 x 500 GB disks will have a capacity of 500 GB.

RAID 5 (Striping with Parity)

With a RAID 5 pool, because data is read from many disks in parallel, as for RAID 0, read performance is good. Write performance is slightly lower because, in addition to writing the data, parity data has to be calculated and written. If a hardware RAID controller is used, this will be done using dedicated hardware; if software RAID is used, the work will be done on the main processor of the storage controller.

The capacity of a RAID 5 pool is reduced by exactly one disks worth of capacity, which is required to store the parity data. For example, a RAID 5 pool made up of 3 x 500 GB disks will have a capacity of 1 TB.

¹Capacity is usually very slightly less because a small but insignificant amount of space is reserved by the RAID controller to store internal metadata.

In principle, a RAID 5 pool could have a very large number of disks. However, the more disks there are, the greater the chance of a double disk failure. If a single disk fails, the data is no longer protected until the disk has been replaced and the pool has been rebuilt by reconstructing all the data from the failed disk and writing it to the new disk. If the disk capacities are very large, it may take many hours to rebuild the pool. If a second disk fails before the rebuild has completed, all the data in the pool will be lost. That is to say, large capacity disks increase the time taken to rebuild the pool, during which time the pool is vulnerable to a second disk failure. Moreover, the chance of a second disk failure increases as the number of disks in the pool increases.

RAID 5 is similar to RAID 0 in that data is striped across multiple disks. However, one disk's worth of space is reserved to store parity data, which can be used to reconstruct the pool in the event of one of its disks failing. With RAID 5, the parity data is distributed across all the disks in the pool. If a single disk fails, each block of data stored on that disk can be reconstructed using the corresponding data block from all the other disks along with the parity block. This means that if a single disk fails, data can still be read, albeit at a rather slower rate (because it needs to be reconstructed, rather than read directly). For this reason, a RAID 5 pool with a disk failure is referred to as a degraded pool.

RAID 6 (Striping with Dual Parity)

RAID 6 is similar to RAID 5 but instead of storing a single disk's worth of parity data, two disk's worth are stored, making the pool capable of withstanding the failure of two disks. However, there is an additional write overhead involved in calculating the double parity data. Since RAID 6 works best with dedicated hardware, RAID 6 is only offered on systems with a hardware RAID controller. Read performance is similar to that of RAID 0 or 5. Since two disks are used for storing parity data, the capacity of a RAID 6 pool made up of 8 x 500 GB disks will be 3 TB.

RAID 10 (Striped Mirroring)

RAID 10 is defined as mirrored stripe sets or also known as RAID 0+1. You can build RAID 10 either directly through the RAID controller (depending on the controller) or by combining software mirroring and controller striping, or vice versa (called RAID 01).

Recurring Snapshot

A snapshot that runs at an administrator-specified time and interval.

Resynchronization

A RAID state that describes the process of integrating a new drive into the RAID.

Round robin DNS

A technique using DNS hostname resolution to balance client connections to multiple servers or cluster nodes. The DNS contains multiple records for the same hostname pointing to the IP addresses of the servers or nodes operating under that name. For lookups of that hostname, the DNS returns a list of all IP addresses with records for the hostname, and rotates the list order for each response. Clients generally use the first entry in the list, so subsequent client connections following DNS lookups for a hostname rotate through servers or cluster nodes.

Router

A router is a device that enables connectivity between Ethernet network segments.

SAN

Short for *Storage Area Network*. Data storage connected to a network that provides network clients access to data using block level protocols. To the clients, the data storage devices appear local rather than remote. An iSCSI SAN is sometimes referred to as an IP-SAN.

SAS

Short for *Serial Attached SCSI*. It is a point-to-point serial protocol that replaces parallel SCSI bus technology (multidrop) and uses the standard SCSI command set. It has no termination issues, supports up to 16,384 devices (using expanders), and eliminates clock skew. It consists of an Initiator that originates device service requests, a Target containing logical units that receives device service requests, and a Service Delivery Subsystem that transmits information between the Initiator and the Target.

Server Number

A numeric derived from the MAC address of the primary Ethernet port that is used to uniquely identify a SnapServer or SnapScale node.

Session

When an initiator wants to establish a connection with a target, it establishes what is known as an iSCSI session. A session consists of one or more TCP/IP connections between an initiator and a target. Sessions are normally established (or re-established) automatically when the host computer starts up, although they also can be established (and broken) manually.

Share

A virtual folder that maps to the root of a volume or a directory on the volume. Permissions are assigned to a share that determine access for specific users and groups.

Share Access

Permissions granted or denied to users and groups that control user and group access to the files.

S.M.A.R.T.

Short for *Self Monitoring, Analysis and Reporting Technology*. A standard mechanism for querying disk drives to monitor performance and reliability attributes, such as temperature, read error rates and seek times. S.M.A.R.T. systems are built into most modern disk drives.

SMB

Short for *Server Message Block*. A protocol for Windows clients. SMB uses the TCP/IP protocol. It is viewed as a complement to the existing Internet application protocols such as FTP and HTTP. With SMB, you can access local server files, obtain read-write privileges to local server files, share files with other clients, and restore connections automatically if the network fails.

SMTP

Short for *Simple Mail Transfer Protocol*. A TCP/IP protocol used for sending and receiving email.

Snap EDR

A SnapExtension that copies the contents of a share from one SnapServer to another share on one or more SnapServers. Snap EDR is designed to work with SnapServers and other SnapServer Storage Solutions.

SnapExtension

An application that extends the functionality of a SnapServer or SnapScale cluster.

SnapServer Manager

The SnapServer Manager (SSM) is a Java-based utility for discovering and monitoring SnapServers and SnapScale clusters.

Snapshot

A method for producing a point-in-time image of a logical drive that results in a consistent, stable, point-in-time image of a volume (filesystem) used for backup purposes. In the process of initiating a snapshot, no data is actually copied from the volume. However as new writes are made to a volume, existing data blocks are copied to the snapshot pool before the new data is written to the volume.

Snapshot Chaining

A native technology in which all snapshots of a volume depend on successive snapshots for part of their content.

Snapshot Pool

Disk space reserved within a RAID for the storage of snapshot data. In the default storage configuration of many SnapServers, twenty percent of the RAID capacity is allocated to the snapshot pool.

Snapshot Share

A virtual folder that allows access to all current snapshots at the same directory level as the original share on which it is based.

SNMP

Short for *Simple Network Management Protocol*. A system to monitor and manage network devices such as computers, routers, bridges, and hubs. SNMP views a network as a collection of cooperating, communicating devices, consisting of managers and agents.

SSH

Short for *Secure Shell*. A service that provides a remote console for special system administration and customer support access to the server. SSH is similar to telnet but more secure, providing strong encryption so that no passwords cross the network in clear text.

SSL

Short for *Secure Sockets Layer*. A protocol for managing the security of a message sent on the Internet. It is a type of technology that provides data encryption, server authentication, message integrity, and client authentication for any TCP/IP connection.

Standalone

A network bonding mode which treats each port as a separate interface. This configuration should be used only in multihomed environments in which network storage resources must reside on two separate subnets.

Static IP Address

An IP address defined by the system administrator rather than by an automated system, such as DHCP.

Storage Area Network

See [SAN](#).

Subnet Mask

A portion of a network that shares a common address component. On TCP/IP networks, subnets are all devices with IP addresses that have the same prefix.

Target

A target is a device (peripheral) that responds to an operation requested by an initiator (host system). Although peripherals are generally targets, a peripheral may be required to act temporarily as an initiator for some commands (for example, SCSI COPY command).

Targets are embedded in iSCSI storage controllers. They are the software that makes the RAID storage available to host computers, making it appear just like any other sort of disk drive.

TCP/IP

Short for *Transmission Control Protocol/Internet Protocol*. The basic protocol used for data transmission over the Internet.

Trap

A signal from a device informing an SNMP management program that an event has occurred.

U

A standard unit of measure for designating the height in computer enclosures and rack cabinets. One U equals 1.75 inches. For example, a 3U server chassis is 5.25 inches high.

UDP

Short for *User Datagram Protocol*. A communications protocol for sending messages between computers in a network that uses the Internet Protocol (IP). UDP is an alternative to the Transmission Control Protocol but, unlike TCP, does not guarantee reliability or ordering of data packets.

UID

Short for *User Identification*. A unique ID assigned to each user on a SnapServer for security purposes.

UNC

Short for *Universal Naming Convention*. In a network, a way to identify a shared file in a computer without having to specify (or know) the storage device it is on. In the Windows OS, the UNC name format is as follows:

```
\\server_name\share_name\path\file_name
```

UPS

Short for *Uninterruptible Power Supply*. A device that allows a computer to keep running for a short time when the primary power source is lost. It also provides protection from power surges. A UPS device contains a battery that starts when the device senses a loss of power from the primary source.

URL

Short for *Uniform Resource Locator*. A Web address.

USB Port

USB is short for *Universal Serial Bus*. A USB port is a hardware interface for low-speed peripherals such as the keyboard, mouse, joystick, scanner, printer, and telephony devices.

VDS

Short for *Virtual Disk Service*. VDS is a feature of Microsoft Windows (from Windows Server 2003 onwards). It provides a consistent interface for managing storage devices and creating volumes. Each vendor of a storage solution can write their own hardware provider module that enables the standard set of VDS commands to be used with different enclosures. Thus, multiple storage systems by different vendors can be controlled using the same set of VDS commands.

Volumes

A logical partition of a RAID's storage space that contains a filesystem. Volumes are created from storage pools, using unused capacity in a pool. Most can be extended in size, so long as there is free capacity in the storage pool.

VSS

Short for *Volume Shadow Copy Service*. A low level communications interface that enables volumes to be backed up without having to halt all applications that are reading or writing the volumes. Microsoft VSS provides a mechanism for creating consistent point-in-time copies of data known as shadow copies.

Web Management Interface

A Web-based utility used for configuration and ongoing maintenance, such as monitoring server conditions, configuring email alerts for key events, or for SNMP management.

Web View

The Web-browser page that opens when users access a SnapServer or SnapScale cluster using their Web browsers, and displays a list of all shares.

Windows Domain Authentication

Windows-based networks use a domain controller to store user credentials. The domain controller can validate all authentication requests on behalf of other systems in the domain. The domain controller can also generate encrypted challenges to test the validity of user credentials. Other systems use encrypted challenges to respond to CIFS/SMB clients that request access to a share.

WINS

Short for *Windows Internet Naming Service*. The server that locates network resources in a TCP/IP-based Windows network by automatically configuring and maintaining the name and IP address mapping tables.

Workgroup

A collection of computers that are grouped for sharing resources such as data and peripherals over a LAN. Each workgroup is identified by a unique name.

Symbols

> (menu flow indicator) **PR-iv**

Numerics

802.3ad link aggregation **4-7**

A

A record (DNS) **4-7**

ACLs

setting file-level permissions (Windows) **B-5**

Active Directory

and name resolution servers **4-11**

joining AD domain **4-15**

Active Directory Service

SnapScale interoperability with **4-12**

Active Users page **7-3**

add new drives **5-46**

adding nodes **5-37**

adjusting snapshot space **5-25**

admin password

changing **9-8**

default **6-2**

Admin password synchronization warning. **6-17**

Administration page **9-3**

ADS **4-12**

Advanced Share Properties **6-7, 6-8**

ALB **4-7**

alert definitions **PR-iv**

alert messages **2-15**

Application Notes **A-3**

Authentication

default settings **6-2**

HTTPS/HTTP **4-23**

Kerberos **4-12**

NIS domain **4-18**

Automatic Load Balancing (ALB) **4-7**

automatic shutdown **3-5**

automatic update checking **8-8**

Average Usage **7-6**

B

backup solutions **A-1**

bond type

change process **4-8**

changing **4-8**

definitions **4-6**

C

CA Unicenter TNg **4-22**

capacity balance **5-44**

Capacity Balancer **1-5**

Capacity Balancer, see *Data Balancer*

change password **9-8**

client access, configuring

FTP **4-19**

HTTPS/HTTP **4-23**

NFS **4-16**

Windows SMB **4-13**

Client network **1-2, 1-5, 4-2**

cluster management name **1-2**

cluster name **1-2**

Cluster Total Usage **7-5**

community strings

read-only **4-22**

read-write **4-22**

connecting

a Mac OS X client **4-13**

a Windows client **4-13**

conventions, typographical **PR-iv**

Current Usage **7-5**

customer support **PR-iii**

D

- Data Balancer 1-5, 5-9
- data import 8-2
- data protection tasks 5-21
- Data Replication Count 1-2, 2-7, 3-2, 5-4, 5-37
- date and time settings 3-3
- Debug Logging 8-11
- defaults
 - admin password 6-2
 - TCP/IP 4-5
- directories, home 6-37
- disabling snapshots on cluster 5-22
- disk drives
 - LED indicators D-1
- DNS A record 4-7
- domain search
 - authentication required 5-19, 6-12, 6-27, 6-30
- domains
 - joining ADS 4-12, 4-15, 6-4
 - joining NIS 4-18
- download website link PR-iii
- drives
 - adding 5-44
 - considerations 5-44
 - failed 5-44
 - hot swap 5-43
 - installing 5-45
 - replacing 5-43

E

- email notification of server events 8-14
- Ethernet, see *Gigabit Ethernet*
- expansion kits 5-37
- exports file, NFS 6-5

F

- failed drive 5-44
- Failover
- failover 1-5
- files, setting permissions for B-4
- FTP
 - connecting via 4-20

G

- GID 6-3
- Groups
 - creating local 6-22
 - file-level access for B-4
 - joining NIS domain 4-18

H

- hardware information pop-up 2-17
- High Water (HW) Marks 7-5
- High Water Reset 7-5
- historical network usage 7-6
- home directories 6-37
- Home page 9-2
- hot spares 5-4
- hot swap drives 5-43
- HP Open View 4-22
- HTTPS/HTTP, configuring 4-23

I

- ID mapping 6-25
- incorporate new drives 5-46
- Individual/Total Network Usage 7-5
- Initial Setup Wizard 2-3, 2-12
- internal temperature, e-mail notification of 8-15
- IQNs for iSCSI disks 5-28
- iSCSI disks 5-27
 - backing up 5-32
 - configuring iSNS 4-27
 - creating 5-29
 - LUNs 5-34
 - multi-initiator support 5-28
 - name (alias) 5-28
- iSNS 4-27

K

- Kerberos 4-12

L

- LEDs
 - disk drive indicators D-1
 - network D-2
 - power/unit status D-2

- system/status **D-2**
- understanding **D-1**
- Link Aggregation (802.3ad) **4-7**
- Load Balance (ALB) **4-7**
- local groups **6-22**

M

- Macintosh, supported OS versions **1-4**
- maintenance
 - data import **8-2**
 - OS update **8-7**
 - shutdown and restart **8-2**
 - support **8-9**
 - tools **8-13**
- Management IP **1-2**
- Management node **1-2**
- manual check for updates **8-9**
- mapping, ID **6-25**
- menu flow indicator **PR-iv**
- message in Web Management Interface **2-15**
- monitor network traffic **7-4**
- monitoring
 - system **7-1**
- multicast **9-4**

N

- network
 - access **4-1**
 - current settings **4-2**
 - LED indicators **D-2**
- network bonding, see *Failover*
- Network Monitor
 - Average Usage **7-6**
 - Current Usage **7-5**
 - historical usage over time **7-6**
- Network Monitor page **7-4**
- Network Time Protocol (NTP) **3-4**
- new drives detected **5-46**
- NFS
 - access **4-16**
 - configuring **4-16**
 - exports file **6-5**
 - share-level permissions **6-13**
- NIS domains **4-18**
- Node Number **7-2**
- Node Properties **5-36**

- nodes
 - adding **5-37**
 - default page **5-35**
 - Properties page for nodes **5-36**

O

- Open Files page **7-4**
- OS update **8-7**
- Overland technical support **PR-iii**

P

- password
 - changing **9-8**
 - default for admin account **6-2**
 - unlock **6-20**
- paths
 - connecting via web browser **4-24**
- peer sets
 - basics **5-4**
 - data recovery **5-2**
 - definition **1-2**
 - formation **5-2**
 - options **5-5**
 - overview **5-1**
 - page overview **5-5**
- permissions
 - share- and file-level interaction **6-11**
 - file-level, default behavior **B-5**
- Phone Home **8-10**
- Phone home support **8-10**
- power/status LED **D-2**
- product documentation **PR-iii**
- proxy server **8-9**

Q

- Quotas
 - defaults **5-15**
 - deleting **5-20**
 - editing **5-20**
 - initial page **5-17**
 - search for usage **5-17**
 - usage calculation **5-16**

R

RAID

- types defined **GL-10**

RAINcloudOS

- specifications **1-3**

- updating **8-8**

RapidRebuild **5-2**

- reboot, setting up alert for **8-15**

- reducing snapshot space **5-22, 5-25**

- registration **8-12**

- remote SnapServer discovery **9-7**

- replacing disks **5-43**

- replacing drives **5-43**

- replication **A-1**

- reset options **D-3**

- restart **8-2**

S

security

- guides **6-3**

- models **6-25**

- shares **6-5**

- Windows ACLs **B-4**

- server registration, online **8-12**

Shares **6-5**

- delete **6-9**

- edit properties **6-8**

- shutdown **8-2**

- Simple Network Management Protocol, see *SNMP*

- site map **2-16, 9-1**

- server links **3-1**

SMB **4-10**

- SMB or NFS backup **A-3**

SMB2 **4-14**Snap EDR **A-1**Snap Finder **9-6**SnapExtensions **9-5**

- snapshot space **5-22**

snapshots

- create **5-24**

- default page **5-22**

- overview **5-22**

- scheduled **5-26**

- shares **B-3**

SNMP **4-22**

- enable traps **4-23**

- read-only **4-22**

- read-write **4-22**

- software update **PR-iii, 8-7**

- Spare Disk Balancer, see *Spare Distributor*

- Spare Disks page **5-6**

- Spare Distributor **1-5, 5-7**

- specifications, RAINcloudOS **1-3**

storage

- Nodes default page **5-35**

- Volumes default page **5-11**

- Storage network **1-2, 1-5, 4-4**

- support **8-9**

- Switch Trunking **4-7**

- synchronize Admin password warning **6-17**

- system monitor **7-1**

- system reset **D-3**

- system/status LED **D-2**

T

TCP/IP

- options **4-5**

- technical support **PR-iii**

- Tivoli NetView **4-22**

- tools **8-13**

Traditional RAID

- quotas **5-15**

- troubleshooting **D-1**

- typographical conventions **PR-iv**

UUID **6-3**

- Uninitialized node **1-2**

- Uninterruptable Power Supplies (UPS) **3-5**

- unlock a user password **6-20**

- updates manual check **8-9**

- updates to RAINcloudOS **8-7**

UPS

- configuring **3-5**

- enabling support for **3-5**

- low-power warning **3-5**

- uptime **9-4**

- user names, active **7-3**

users

- creating local **6-16**

- file-level access for **B-4**

- Utility IP address **4-9, A-3**

V

- VDS-based iSCSI disks **5-29**
- volumes
 - capacity reached alert **8-15**
 - default page **5-11**
 - expanding capacity of **5-14**
 - Volume Properties **5-14**
- VSS-based iSCSI disks **5-29**

W

- warranty activation **8-12**
- Web Management Interface
 - alert messages **2-15**
- Web Management Interface, overview **2-13**
- Web Root **4-24**
- Web Server **4-24**
- Windows
 - connecting from a client **4-13**
 - enabling guest account access **4-14, 4-16**
 - guest account access **4-13**
 - name resolution server support **4-11**
 - networking (SMB) **4-10**
 - security, joining
 - active directory domain **4-15**
 - see also *Active Directory*
 - see also *Authentication*
- Windows Active Directory
 - setup **4-12, 4-15, 6-4**
 - Shares **6-5**
- workgroup environment **4-12**
- workgroup, joining **4-13**