Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x

F13769-27
June 2025

Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x,

F13769-27

Primary Author:

Contents

2 Configuring the Appliance

10 Integration

1 About Oracle ZFS Storage Appliance

The Oracle ZFS Storage Appliance (appliance) family of products provides efficient file and block data services to clients over a network, and a rich set of data services that can be applied to the data stored on the system.

For information about configuring and working with the Oracle ZFS Storage Appliance product, see the following sections:

- Oracle ZFS Storage Appliance Key Features
- Accessibility, Diversity, and Inclusion
- Supported Protocols
- Oracle ZFS Storage Appliance Data Services
- Data Availability
- Browser User Interface (BUI)
- Network Icons
- Dashboard Icons
- Analytics Icons
- Identity Mapping Icons
- Supported Browsers
- Command Line Interface (CLI)
- Working with CLI Scripting

Oracle ZFS Storage Appliance Key Features

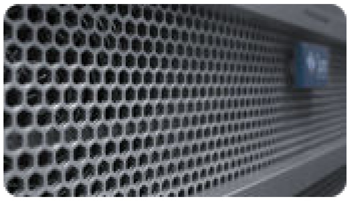

Oracle ZFS Storage Appliance includes technologies to deliver the best storage price/performance and unprecedented observability of your workloads in production, including:

- DTrace Analytics, a system for dynamically observing the behavior of your system in real-time and viewing data graphically
- The ZFS Hybrid Storage Pool, composed of optional Flash-memory devices for acceleration of reads and writes, low-power, high-capacity disks, and DRAM memory, all managed transparently as a single data hierarchy
- Support for a variety of hardware

For more information about analytics and hardware, refer to the documentation in the Oracle Help Center (https://docs.oracle.com/en/storage/).

Accessibility, Diversity, and Inclusion

This section documents the accessibility features of Oracle ZFS Storage Appliance and also Oracle's commitment to using language that is ethnically and culturally diverse and inclusive.

Accessibility

The Oracle ZFS Storage Appliance command-line interface (CLI) is an alternative, and equivalent, way to access the browser user interface (BUI) features and functionality. Because the operating systems that run on Oracle ZFS Storage Appliance systems support assistive technologies to read the content of the screen, you can use the CLI as an equivalent means to access the color-based, mouse-based, and other visual-based utilities that are part of the BUI. For example, you can use a keyboard to enter CLI commands to identify faulted hardware components, check system status, and monitor system health.

Oracle strives to make its products, services, and supporting documentation usable and accessible to the disabled community. To that end, products, services, and documentation include features that make the product accessible to users of assistive technology. For more information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers and partners we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

Supported Protocols

Oracle ZFS Storage Appliance supports a variety of industry-standard client protocols, including NFS, iSCSI, SMB, FTP, HTTP, NDMP, Fibre Channel, SRP, iSER, and SFTP.

For information on these protocols, see the following sections:

- SAN Fibre Channel Configuration
- Configuring SAN iSER Targets
- NFS Configuration
- iSCSI Configuration
- SMB Configuration
- FTP Configuration
- HTTP Configuration
- NDMP Configuration
- SFTP Configuration

- SRP Configuration

Appliance Data Services

To manage the data that you export using these protocols, you can configure Oracle ZFS Storage Appliance using the built-in collection of advanced data services, including:

- RAID-Z (RAID-5 and RAID-6), mirrored, and striped disk configurations: See Configuring Storage
- Unlimited read-only and read-write snapshots, with snapshot schedules: See Snapshots and Clones
- Controlling the elimination of duplicate copies of data: See Data Deduplication
- Built-in data compression: See Data Compression
- Remote replication of data for disaster recovery: See Remote Replication
- Active-active clustering for high availability: See Appliance Cluster Configuration
- Thin provisioning of iSCSI LUNs: See iSCSI Configuration
- Virus scanning and quarantine: See Virus Scan Configuration
- NDMP backup and restore: See NDMP Configuration

> **Note:**
>
> Replication and Cloning are licensed features for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Data Availability

To maximize the availability of your data in production, Oracle ZFS Storag Appliance includes a complete end-to-end architecture for data integrity, including redundancies at every level of the stack. Key features include:

- Predictive self-healing and diagnosis of all system hardware failures: CPUs, DRAM, I/O cards, disks, fans, power supplies
- ZFS end-to-end data checksums of all data and metadata, protecting data throughout the stack
- RAID-6 (double- and triple-parity) and optional RAID-6 across disk shelves
- Active-active clustering for high availability: See Appliance Cluster Configuration
- Link aggregations and IP network multipathing for network failure protection: See Network Configuration
- I/O Multipathing between the controller and disk shelves
- Integrated software restart of all system software services: See Appliance Services
- Phone Home of telemetry for all software and hardware issues: See Phone Home Configuration
- Lights-out Management of each system for remote power control and console access

Browser User Interface (BUI)

The Oracle ZFS Storage Appliance Browser User Interface (BUI) is the graphical tool for administration of the appliance. The BUI provides an intuitive environment for administration tasks, visualizing concepts, and analyzing performance data. The BUI provides an uncluttered environment for visualizing system behavior and identifying performance issues with the appliance.

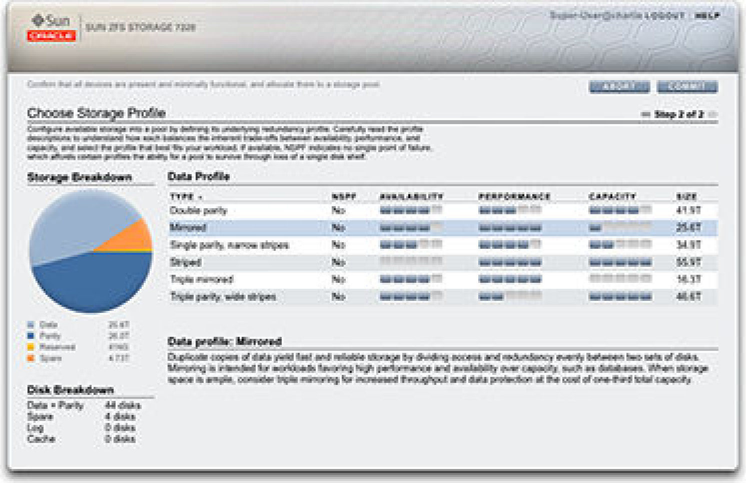

Confirm that all devices are present and minimally functional, and allocate them to a storage pool.

Choose Storage Profile

Configure available storage into a pool by defining its underlying redundancy profile. Carefully read the profile descriptions to understand how each balances the inherent trade-offs between availability, performance, and capacity, and select the profile that best fits your workload. If available, NSPF indicates no single point of failure, which affords certain profiles the ability for a pool to survive through loss of a single disk shelf.

Storage Breakdown

🟦 Data	25.6T
🟦 Parity	26.0T
🟨 Reserved	419G
🟧 Spare	4.73T

Disk Breakdown

Data + Parity	44 disks
Spare	4 disks
Log	0 disks
Cache	0 disks

Data Profile

TYPE ▴	NSPF	AVAILABILITY	PERFORMANCE	CAPACITY	SIZE
Double parity	No	■■■■■□	■■■■□□	■■■■■□	41.5T
Mirrored	No	■■■■■□	■■■■■■	■□□□□□	25.6T
Single parity, narrow stripes	No	■■■■□□	■■■■■□	■■□□□□	34.9T
Striped	No	□□□□□□	■■■■■■	■■■■■■	55.9T
Triple mirrored	No	■■■■■□	■■■■■■	□□□□□□	16.3T
Triple parity, wide stripes	No	■■■■■□	■■□□□□	■■■■■■	46.6T

Data profile: Mirrored

Duplicate copies of data yield fast and reliable storage by dividing access and redundancy evenly between two sets of disks. Mirroring is intended for workloads favoring high performance and availability over capacity, such as databases. When storage space is ample, consider triple mirroring for increased throughput and data protection at the cost of one-third total capacity.

Direct your browser to the system by using either the *IP address* or *host name* you assigned to the `NET-0` port during initial configuration as follows: `https://ipaddress:215` or `https://hostname:215`

The login screen appears. If multi-factor authentication is enabled, you might be prompted for additional information during the login sequence.

The online help link in the top-right of the BUI screen is context-sensitive. For every top-level and second-level screen in the BUI, the associated help page appears when you click the **Help** button.

To view the online help library, which contains all online help documents, click the **Help** button to go to online help, and click the **Release OS8.8.x** breadcrumb. To view the documents in a different language, use the language selector, and then select a document to view its latest translation.

The masthead contains several interface elements for navigation and notification, as well as primary functionality. At left, from top to bottom, are the Sun/Oracle logo, a hardware **model badge**, and hardware **power off/restart** button. Across the right, again from top to bottom: **login identification**, **logout**, **help, main navigation**, and **subnavigation**.

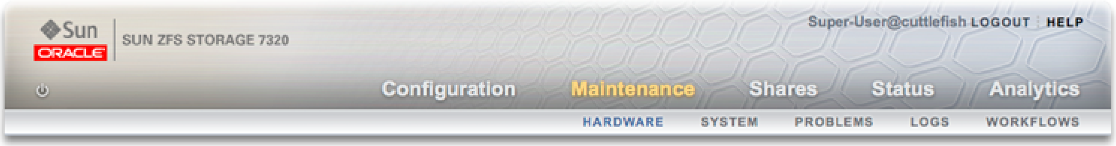

System alerts appear in the Masthead as they are triggered. If multiple alerts are triggered sequentially, refer to the list of recent alerts found on the **Dashboard** screen or the full log available on the **Logs** screen.

Use the main navigation links to view between the **Configuration**, **Maintenance**, **Shares**, **Status**, and **Analytics** areas of the BUI. Use sub-navigation links to access features and functions within each area.

If you provide a session annotation, it appears beneath your **login ID** and the **logout** control. To change your session annotation for subsequent administrative actions without logging out, click on the text link. For details about session annotations, see Configuring Users.

The title bar appears below the Masthead and provides local navigation and functions that vary depending on the current view.

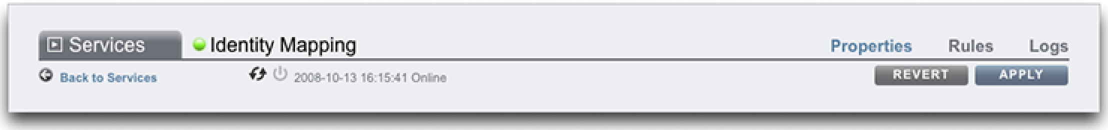

For example, the **Identity mapping** service title bar enables the following:

- Navigation to the full list of services through the side panel
- Controls to enable or disable the **Identity Mapping** service
- A view of **Identity Mapping** uptime
- Navigation to the **Properties**, **Rules** and **Logs** screens for your **Identity Mapping** service
- Button to **Apply** configuration changes made on the current screen
- Button to **Revert** configuration changes applied on the current screen

To quickly navigate between **Service** and **Project** views, open and close the side panel by clicking the title or the reveal arrow.

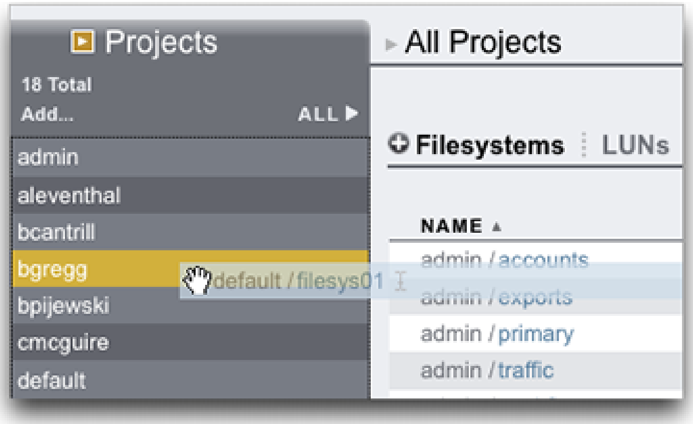

To add projects, click the **Add** link in the sidebar.

To move shares between projects, click the move icon ✥ and drag a filesystem share to the appropriate project in the side panel. Note that dragging a share into another project will change its properties if they are set to be inherited from its parent project.

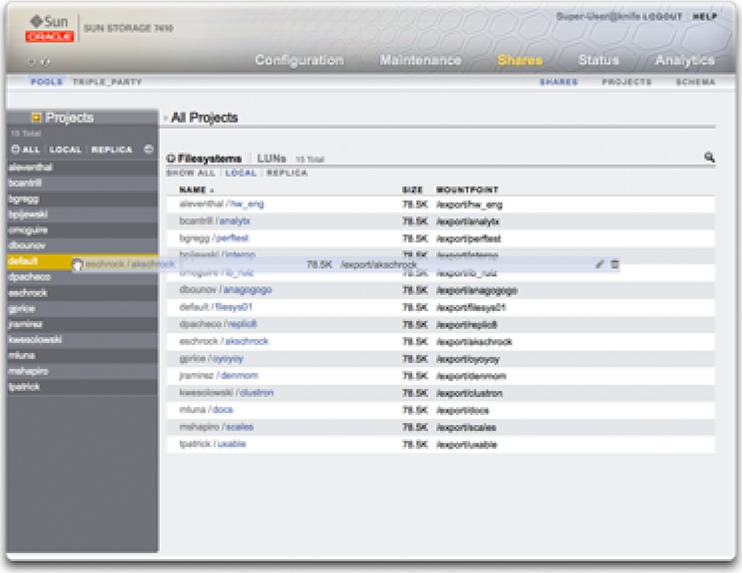

NAME	SIZE	MOUNTPOINT
aleventhal / hw_eng	78.5K	/export/hw_eng
bcantrill / analytx	78.5K	/export/analytx
bgregg / perftest	78.5K	/export/perftest
bpijewski / interno	78.5K	/export/interno
eschrock / akschrock	78.5K	/export/akschrock
cmcguire / lb_rulz	78.5K	/export/lb_rulz
dbounov / anagogogo	78.5K	/export/anagogogo
default / filesys01	78.5K	/export/filesys01
dpacheco / replic8	78.5K	/export/replic8
eschrock / akschrock	78.5K	/export/akschrock
gprice / oyoyoy	78.5K	/export/oyoyoy
jramirez / denmom	78.5K	/export/denmom
kwesolowski / clustron	78.5K	/export/clustron
mluna / docs	78.5K	/export/docs
mshapiro / scales	78.5K	/export/scales
tpatrick / usable	78.5K	/export/usable

Most BUI controls use standard web form inputs; however, there are a few key exceptions worth noting.

Table 1-1 Key Web Form Exceptions

Summary of BUI Controls	Description
Modify a property	Click the edit icon ✎ and complete the dialog box
Add a list item or property entry	Click the add icon ⊕
Remove a list item or property entry	Click the remove icon ⊖
Save changes	Click the **Apply** button
Undo saved changes	Click the **Revert** button
Delete an item from a list	Click the trash icon 🗑 (hover the mouse over the item row to see the icon)
Search for an item in a list	Click the search icon 🔍 at the top right of the list
Sort by list headings	Click on the bold sub-headings to re-sort the list
Move or drag an item	Click the move icon ✥
Rename an item	Click the rename icon I
View details about your system	Oracle logo. To navigate to the Oracle product page for your model, click the model badge.
Automatically open side panel	Drag an item to the side panel

When setting permissions, the RWX boxes are clickable targets. Clicking on the access group label (**User**, **Group**, **Other**) toggles all permissions for that label on and off.

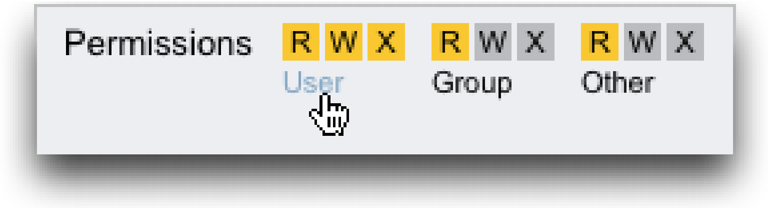

To edit **Share** properties, deselect the **Inherit from project** check box.

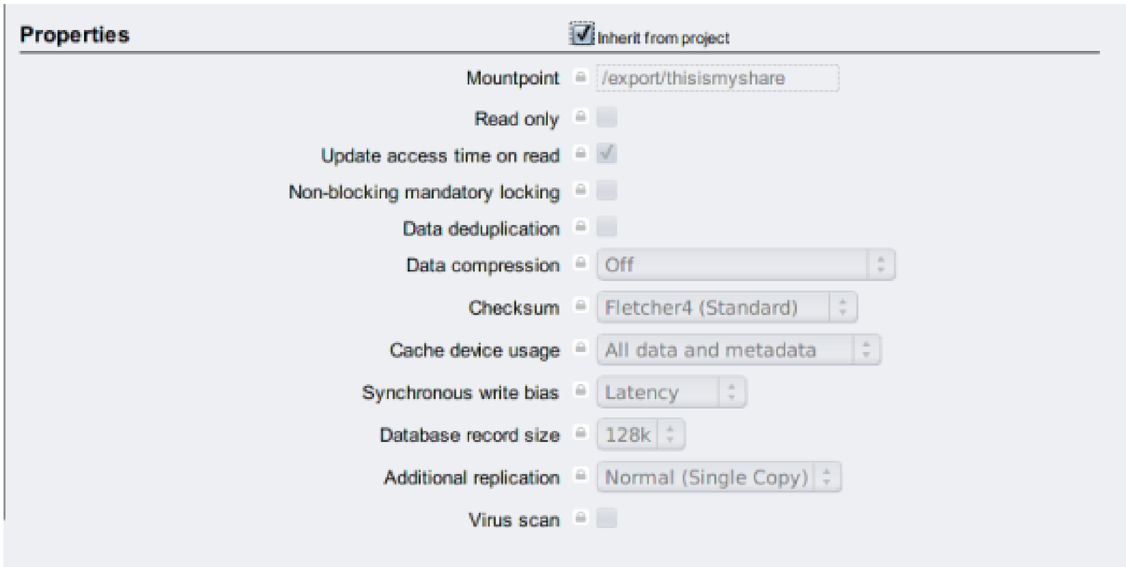

Properties ☑ Inherit from project

Mountpoint: /export/thisismyshare
Read only: ☐
Update access time on read: ☑
Non-blocking mandatory locking: ☐
Data deduplication: ☐
Data compression: Off
Checksum: Fletcher4 (Standard)
Cache device usage: All data and metadata
Synchronous write bias: Latency
Database record size: 128k
Additional replication: Normal (Single Copy)
Virus scan: ☐

To view controls for an item in a list, hover the mouse over the row.

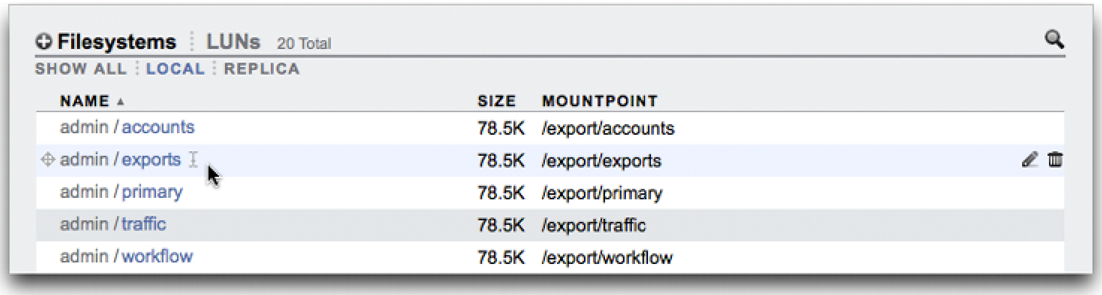

Filesystems | LUNs 20 Total

SHOW ALL | LOCAL | REPLICA

NAME ▲	SIZE	MOUNTPOINT
admin / accounts	78.5K	/export/accounts
admin / exports	78.5K	/export/exports
admin / primary	78.5K	/export/primary
admin / traffic	78.5K	/export/traffic
admin / workflow	78.5K	/export/workflow

All modal dialog boxes have titles and buttons that identify and commit or cancel the current action at top, and content below. The modal content area follows the same interface conventions as the main content area, but are different in that they must be dismissed using the buttons in the title bar before other actions can be performed.

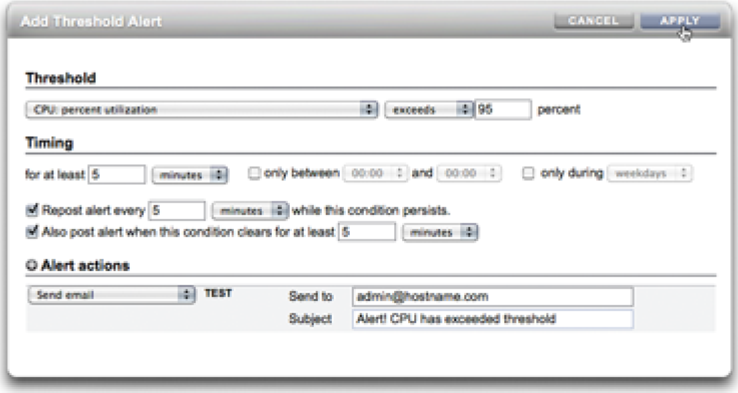

Threshold

CPU: percent utilization | exceeds | 95 percent

Timing

for at least 5 minutes ☐ only between 00:00 and 00:00 ☐ only during weekdays

☑ Repost alert every 5 minutes while this condition persists.
☑ Also post alert when this condition clears for at least 5 minutes

Alert actions

Send email TEST Send to admin@hostname.com
Subject Alert! CPU has exceeded threshold

Icons indicate system status and provide access to functionality, and in most cases serve as buttons to perform actions when clicked. It is useful to hover your mouse over interface icons to view the tool tip. The tables below provide a key to the conventions of the user interface.

The status lights are basic indicators of system health and service state.

Table 1-2 Status Indicators

Icon	Description	Icon	Description
	On	●	Warning
●	Off	●	Disabled

The following icons are found throughout the user interface and cover most of the basic functionality.

Table 1-3 BUI Icons

Icon	Description	Icon	Description
⌶	Rename (edit text)	⇤	Sever
✥	Move	⊞	Clone
Enabled: ✎ Disabled: ✎	Edit	↺	Rollback
Enabled: 🗑 Disabled: 🗑	Destroy	⏻	Appliance power
Enabled: ⊕ Disabled: ⊕	Add	➔	Apply
Enabled: ⊖ Disabled: ⊖	Remove	←	Revert
Enabled: ⊗ Disabled: ⊗	Cancel or close	ⓘ	Info
!	Error	▼	Sort list column (down)
⚠	Alert	▲	Sort list column (up)

Table 1-3 (Cont.) BUI Icons

Icon	Description	Icon	Description
Enabled: ⏻ Disabled: ⏻	On/Off toggle	Enabled: ⏮ Disabled: ⏮	First page
Enabled: ⟳ Disabled: ⟳	Restart	Enabled: ⏪ Disabled: ⏪	Previous page
☼	Locate	Enabled: ⏩ Disabled: ⏩	Next page
Enabled: ⊘ Disabled: ⊘	Disable/Offline	Enabled: ⏭ Disabled: ⏭	Last page
Enabled: 🔒 Disabled: 🔒	Lock	🔍	Search
⁘	Wait spinner	Enabled: ▨ Disabled: ▨	Menu
↻	Reverse direction	Enabled: ▨ Disabled: ▨	Panel

The following icons are used to distinguish different types of objects and provide information of secondary importance.

Table 1-4 Miscellaneous Icons

Icon	Description	Icon	Description
⇶	SAS	⬡	Storage pool
▬	SAS port		

Network Icons

Network icons indicate the state of network devices and type of network datalinks.

Table 1-5 Network Icons

Icon	Description	Icon	Description
▬	Active network device	▨	Active InfiniBand port
▬	Inactive network device	▨	Inactive InfiniBand port
‹··›	Network datalink	⚯	Network datalink for an InfiniBand partition

Table 1-5 (Cont.) Network Icons

Icon	Description	Icon	Description
{∷}	Network datalink aggregation		
‹∘∘∘›	Network datalink VLAN		
{⁘}	Network datalink aggregation VLAN		

Dashboard Icons

The following icons indicate the current state of monitored statistics with respect to user-configurable thresholds set from within **Settings**.

Table 1-6 Dashboard Icons

Icon	Description	Icon	Description
☀	Sunny	🌀	Hurricane
⛅	Partly cloudy	🌀2	Hurricane class 2
☁	Cloudy	🌀3	Hurricane class 3
💧	Rainy	🌀4	Hurricane class 4
⚡	Stormy	🌀5	Hurricane class 5

Analytics Toolbar Icons

This set of icons is used in a toolbar to manipulate display of information within Analytics worksheets.

Table 1-7 Analytics Toolbar Icons

Icon	Description	Icon	Description
⬅	Back	∨	Show minimum
➡	Forward	∧	Show maximum
⏭	Forward to now	〰	Show line graph
⏸	Pause	▲	Show mountain graph
🔍−	Zoom out	⌐	Crop outliers
🔍+	Zoom in	⤡	Sync worksheet to this statistic
◷	Show one minute	⤢	Unsync worksheet statistics

Table 1-7 (Cont.) Analytics Toolbar Icons

Icon	Description	Icon	Description
🕒	Show one hour	🔧	Drilldown
📅	Show one day	⤓	Export statistical data (download to client)
📅	Show one week	●	Save statistical data
📅	Show one month	●	Archive dataset
⤒	Send worksheet with support bundle		

For more information about Analytics, refer to the documentation in the Oracle Help Center (https://docs.oracle.com/en/storage/).

Identity Mapping Icons

These icons indicate the type of role being applied when mapping users and groups between Windows and UNIX.

Table 1-8 Identity Mapping Icons

Icon	Description	Icon	Description
o> o>	Allow Windows to UNIX	<o <o	Allow UNIX to Windows
o-| o-|	Deny Windows to UNIX	|-o |-o	Deny UNIX to Windows
⇔ ⇔	Allow bidirectional		

For more information about Analytics, refer to the documentation on the Oracle Help Center (https://docs.oracle.com/en/storage/).

Related Topics

- Understanding the Appliance Status
- Network Configuration
- Configuring Storage
- Configuring Alerts
- Appliance Services

Supported Browsers

This section defines BUI browser support.

The BUI is fully featured and functional on the following browsers:

- Firefox 10 & newer
- Internet Explorer 11 & newer
- Google Chrome 31 & newer

- Safari 5 & newer
- Edge 98 & newer

BUI elements may be cosmetically imperfect on the following browsers, and some functionality may not be available, although all necessary features function correctly. A warning message appears during log in if you are using one of the following browsers:

- Firefox 6 to 9
- Internet Explorer 11
- Google Chrome 21 to 30
- Safari: NA
- Edge: NA

The following browsers are incompatible, unsupported, and known to have issues; log in will not complete.

- Firefox 5 & older
- Internet Explorer 10 & older
- Google Chrome 20 & older
- Safari 4 & older
- Edge: NA

Related Topics

- Configuring Users
- Setting Appliance Preferences

Command Line Interface (CLI)

The CLI is designed to imitate the capabilities of the BUI, while also providing a powerful scripting environment for performing repetitive tasks. The command line is an efficient and powerful tool for repetitive administrative tasks. The appliance presents a CLI available through either the Console or SSH. There are several situations in which the preferred interaction with the system is the CLI:

- **Network unavailability** - If the network is unavailable, browser-based management is impossible; the only vector for management is the Console, which can only accommodate a text-based interface
- **Expediency** - Starting a browser may be prohibitively time-consuming, especially if you only want to examine a particular aspect of the system or make a quick configuration change
- **Precision** - In some situations, the information provided by the browser may be more qualitative than quantitative in nature, and you need a more precise answer
- **Automation** - Browser-based interaction cannot be easily automated; if you have repetitive or rigidly defined tasks, script the tasks
- **Accessibility** - The CLI is an alternative, and equivalent, way to access the BUI features and functionality. Because the operating systems that run on Oracle ZFS Storage Appliance systems support assistive technologies to read the content of the screen, you can use the CLI as an equivalent means to access the color-based, mouse-based, and other visual-based utilities that are part of the BUI. For example, you can use a keyboard

to enter CLI commands to identify faulted hardware components, check system status, and monitor system health.

Oracle strives to make its products, services, and supporting documentation usable and accessible to the disabled community. To that end, products, services, and documentation include features that make the product accessible to users of assistive technology. For more information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website (http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc).

When navigating through the CLI, there are two principles to be aware of:

- **Tab completion is used extensively** - if you are not sure what to type in any given context, pressing the Tab key will provide you with possible options. Throughout the documentation, pressing Tab is presented as the word "tab" in bold italics.
- **Help is always available** - the `help` command provides context-specific help. Help on a particular topic is available by specifying the topic as an argument to `help`, for example `help commands`. Available topics are displayed by tab-completing the `help` command, or by typing `help topics`.

You can combine these two principles as follows:

```
hostname:> help tab

builtins    commands    general     help        properties  script
```

To log in remotely using the CLI, use an `ssh` client. If you have not followed the instructions in Configuring Users to administer the appliance, you will need to log in as `root`. When you log in, the CLI will present you with a prompt that consists of the hostname, followed by a colon, followed by a greater-than sign:

```
$ ssh root@hostname
Password:
Last login: Fri Oct 14 15:43:05 2022 from example.us.sample
hostname:>
```

If multi-factor authentication is enabled, you might be prompted for additional information during the login sequence.

Related Topics

- Browser User Interface (BUI)
- CLI Contexts
- CLI Properties

CLI Contexts

A central principle in the CLI is the *context* in which commands are executed. The context dictates which elements of the system can be managed and which commands are available. Contexts have a tree structure in which contexts may themselves contain nested contexts and the structure generally mirrors that of the views in the BUI.

The initial context upon login is the *root context*, and serves as the parent or ancestor of all contexts. To navigate to a context, execute the name of the context as a command. For example, the functionality available in the **Configuration** view in the browser is available in the `configuration` context of the CLI. From the root context, this can be accessed by typing it directly:

```
hostname:> configuration
hostname:configuration>
```

Note that the prompt changes to reflect the context, with the context provided between the colon and the greater-than sign in the prompt.

The `show` command shows child contexts. For example, from the `configuration` context:

```
hostname:configuration> show
Children:
                             net => Configure networking
                        services => Configure services
                         version => Display system version
                           users => Configure administrative users
                           roles => Configure administrative roles
                     preferences => Configure user preferences
                          alerts => Configure alerts
                         storage => Configure Storage
```

These child contexts correspond to the views available under the **Configuration** view in the browser, including **Network**, **Services**, **Users**, **Preferences**, and so on. To select one of these child contexts, type its name:

```
hostname:configuration> preferences
hostname:configuration preferences>
```

Navigate to a descendant context directly from an ancestor by specifying the intermediate contexts separated with spaces. For example, to navigate directly to `configuration preferences` from the root context, simply type it:

```
hostname:> configuration preferences
hostname:configuration preferences>
```

Some child contexts are *dynamic* in that they correspond not to fixed views in the browser, but rather to dynamic entities that have been created by either the user or the system. There are two ways to navigate to these contexts: You can use the `select` command followed by the name of the dynamic context, or surround the name of the dynamic context with double quotes. The names of the dynamic contexts contained within a given context are shown using the `list` command. For example, the `users` context is a static context, but each user is its own dynamic context.

```
hostname:> configuration users
hostname:configuration users> list
NAME                     USERNAME                 UID        TYPE
Pat M Doe                pmd                      12345      Dir
Super-User               root                     0          Loc
```

To select the user named `pmd`, issue the command `select pmd` or `"pmd"`:

```
hostname:configuration users> "pmd"
hostname:configuration users pmd
```

Alternately, double quotes, `select` and `destroy` can in some contexts be used to select an entity based on its properties. For example, one could select log entries issued by the `reboot` module in the `maintenance logs system` context by issuing the following command:

```
hostname:maintenance logs system> select module=reboot
hostname:maintenance logs system entry-034> show
Properties:
  timestamp = 2022-10-14 06:24:41
```

```
   module = reboot
 priority = crit
     text = initiated by root on /dev/console syslogd: going down on signal 15
```

As with other commands, `select` or double quotes may be appended to a context-changing command. For example, to select the user named `pmd` from the root context:

```
hostname:> configuration users select pmd
hostname:configuration users pmd>
```

Use the `last` command to navigate to a previously selected or created context. The following example creates a replication action, and then uses the `last` and `get id` commands to retrieve the replication action ID. Then a different action is selected, and the `last` and `get id` commands are used to retrieve the ID of the last-visited replication action.

Using `last`, you can return to the last-visited node:

```
hostname:configuration net interfaces> "igb4"
hostname:configuration net interfaces igb4> done
hostname:configuration net interfaces> last
hostname:configuration net interfaces igb4>
```

The `last` command is also useful to retrieve values that have been automatically set by the appliance during the creation of a dynamic node. For example, each replication action is assigned an ID by the appliance when it is created. Using the `last` command with the `get id` command, you can retrieve the ID without using the name of the replication action:

```
hostname:shares p1/share replication> create
hostname:shares p1/share action (uncommitted)> set target=hostname
                               target = hostname (uncommitted)
hostname:shares p1/share action (uncommitted)> set pool=p0
                               pool = p0 (uncommitted)
hostname:shares p1/share action (uncommitted)> commit
hostname:shares p1/share replication> last get id
                               id = 7034367a-d4d8-e26f-fa93-c3b454e3b595
hostname:shares p1/share replication>
```

Note that when `last` is combined with another command (in this case, `get id`), the command is run in the context of the last-visited node, but the current node remains unchanged.

Because `last` allows you to retrieve the last-visited node and its values without specifying the name of the node, this command is particularly convenient for scripting:

```
script
         project = 'myproj';
         target = 'mytarget';
         target_pool = 'notmypool';

         run('cd /');
         run('shares select ' + project);
         run('replication');
         run('create');
         set('target', target);
         set('pool', target_pool);
         run('commit');
         run('last');
         id = get('id');
         printf("Sending update for replication action id %s ...", id);
         run('sendupdate');
         while (get('state') != 'idle') {
                printf(".");
                run('sleep 1');
        }
        printf("done\n");
 .
```

To return to the previous context, use the `done` command:

```
hostname:configuration> done
hostname:>
```

This returns to the previous context, which is not necessarily the parent context, as follows:

```
hostname:> configuration users select pmd
hostname:configuration users pmd> done
hostname:>
```

The `done` command can be used multiple times to backtrack to earlier contexts:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> select pmd
hostname:configuration users pmd> done
hostname:configuration users> done
hostname:configuration> done
hostname:>
```

To navigate to a parent context, use the `cd` command. Inspired by the classic UNIX command, `cd` takes an argument of "`..`" to denote moving to the parent context:

```
hostname:> configuration users select pmd
hostname:configuration users pmd> cd ..
hostname:configuration users>
```

And as with the UNIX command, "`cd /`" moves to the root context:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> select pmd
hostname:configuration users pmd> cd /
hostname:>
```

And as with its UNIX analogue, "`cd ../..`" may be used to navigate to the grandparent context:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> select pmd
hostname:configuration users pmd> cd ../..
hostname:configuration>
```

Note that the `cd /` and `cd ..` commands support limited variations. For more versatility, use the `top` command and the `up` command.

Use the `top` command to navigate to the root context:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> select pmd
```

```
hostname:configuration users pmd> top
hostname:>
```

Use the `top` command followed by a context name to directly navigate to the specified context relative to the root context. For example, to directly navigate from context `configuration users` to context `configuration services`, use the `top configuration services` command:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> top configuration services
hostname:configuration services>
```

When the `top` command is used in conjunction with a specific context, the `done` command can be used to navigate back to the context before the `top` command was executed. In the following example, the first `done` command returns to the previous context. The second `done` command returns to the context before the `top` command. The third `done` command returns to the context two nodes before the `top` command.

```
hostname:> maintenance system
hostname:maintenance system> updates
hostname:maintenance system updates> top configuration services
hostname:configuration services> ftp
hostname:configuration services ftp> done
hostname:configuration services> done
hostname:maintenance system updates> done
hostname:>
```

Like the `cd ..` command, the `up` command can be used to navigate to the parent context:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> select pmd
hostname:configuration users pmd> up
hostname:configuration users>
```

Additionally, you can go to a context *n* nodes up from the current context by repeating the `up` command *n* times:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> select pmd
hostname:configuration users pmd> up up
hostname:configuration>
```

To go back to a specific context relative to the current parent context, enter the context name after the `up` command. Likewise, use the `up up` command followed by a context name to go back to a specific context relative to the current grandparent context. For example, to go from context `configuration users pmd` to context `configuration services`, use the command `up up services`:

```
hostname:> configuration
hostname:configuration> users
hostname:configuration users> select pmd
hostname:configuration users pmd> up up services
hostname:configuration services>
```

When the `up` command is used in conjunction with a specific context, the `done` command can be used to navigate back to the context before the `up` command was executed. In the following example, the first `done` command returns to the context before the `up` command. The second `done` command returns to the context two nodes before the `up` command, and the third `done` command returns to the context three nodes before the `up` command.

```
hostname:> configuration
hostname:configuration> services
hostname:configuration services> ftp
hostname:configuration services ftp> up http
hostname:configuration services http> done
hostname:configuration services ftp> done
hostname:configuration services> done
hostname:configuration> done
hostname:>
```

Context names will tab complete, be they static contexts (via normal command completion) or dynamic contexts (via command completion of the `select` command). Following is an example of selecting the user named `pmd` from the root context with just fifteen keystrokes, instead of the thirty-one that would be required without tab completion:

```
hostname:> configtab

hostname:> configuration utab

hostname:> configuration users setab

hostname:> configuration users select tab

pmd   root
hostname:> configuration users select ptab

hostname:> configuration users select pmdenter

hostname:configuration users pmd>
```

Once in a context, execute context-specific commands. For example, to get the current user's preferences, execute the `get` command from the `configuration preferences` context:

```
hostname:configuration preferences> get
                      locale = C
                login_screen = status/dashboard
             session_timeout = 15
          session_annotation =
          advanced_analytics = false
```

If there is input following a command that changes context, that command will be executed in the target context, but control will return to the calling context. For example, to get preferences from the root context without changing context, append the `get` command to the context navigation commands:

```
hostname:> configuration preferences get
                      locale = C
                login_screen = status/dashboard
             session_timeout = 15
          session_annotation =
          advanced_analytics = false
```

When creating a new entity in the system, the context associated with the new entity will often be created in an *uncommitted* state. For example, create a threshold alert by executing the `create` command from the `configuration alerts threshold` context:

```
hostname:> configuration alerts thresholds create
hostname:configuration alerts threshold (uncommitted)>
```

The `(uncommitted)` in the prompt denotes that this an uncommitted context. An uncommitted entity is committed via the `commit` command; any attempt to navigate away from the uncommitted context will prompt for confirmation:

```
hostname:configuration alerts threshold (uncommitted)> cd /
Leaving will abort creation of "threshold". Are you sure? (Y/N)
```

When committing an uncommitted entity, the properties associated with the new entity will be validated, and an error will be generated if the entity cannot be created. For example, the creation of a new threshold alert requires the specification of a statistic name; failure to set this name results in an error:

```
hostname:configuration alerts threshold (uncommitted)> commit
error: missing value for property "statname"
```

To resolve the problem, address the error and reattempt the commit:

```
hostname:configuration alerts threshold (uncommitted)> set statname=cpu.utilization
                statname = cpu.utilization (uncommitted)
hostname:configuration alerts threshold (uncommitted)> commit
error: missing value for property "limit"
hostname:configuration alerts threshold (uncommitted)> set limit=90
                   limit = 90 (uncommitted)
hostname:configuration alerts threshold (uncommitted)> commit
hostname:configuration alerts thresholds> list
THRESHOLD          LIMIT       TYPE STATNAME
threshold-000         90     normal cpu.utilization
```

Related Topics

- Command Line Interface (CLI)
- CLI Properties

CLI Properties

Properties are typed name/value pairs that are associated with a context. Properties for a given context can be ascertained by running the `help properties` command. Following is an example of retrieving the properties associated with a user's preferences:

```
hostname:configuration preferences> help properties
Properties that are valid in this context:

locale             => Locality

login_screen       => Initial login screen

session_timeout    => Session timeout

session_annotation => Current session annotation

advanced_analytics => Make available advanced analytics statistics
```

The properties of a given context can be retrieved with the `get` command. Following is an example of using the `get` command to retrieve a user's preferences:

```
hostname:configuration preferences> get
                       locale = C
                 login_screen = status/dashboard
              session_timeout = 15
           session_annotation =
           advanced_analytics = false
```

The `get` command will return any properties provided to it as arguments. For example, to get the value of the `login_screen` property:

```
hostname:configuration preferences> get login_screen
                 login_screen = status/dashboard
```

The `get` command will tab complete with the names of the available properties. For example, to see a list of available properties for the iSCSI service:

```
hostname:> configuration services iscsi get tab

<status>            isns_server         radius_secret       target_chap_name
isns_access         radius_access       radius_server       target_chap_secret
```

The `select` command, or a command surrounded by double quotes, will select a dynamic node by property. For example, to select `key-000` by user:

```
hostname:configuration services sftp keys> show
Keys:

NAME    MODIFIED            CIPHER   USER    COMMENT
key-000 2022-6-5 19:48:23   RSA      u1      1

hostname:configuration services sftp keys> "user=u1"
hostname:configuration services sftp key-000>
```

The `set` command will set a property to a specified value, with the property name and its value separated by an equals sign. For example, to set the `login_screen` property to be "`shares`":

```
hostname:configuration preferences> set login_screen=shares
                 login_screen = shares (uncommitted)
```

Note that in the case of properties that constitute state on the appliance, setting the property does *not* change the value, but rather records the set value and indicates that the value of the property is uncommitted.

To force set property values to take effect, they must be explicitly committed, allowing multiple values to be changed as a single, coherent change. To commit any uncommitted property values, use the `commit` command:

```
hostname:configuration preferences> get login_screen
                 login_screen = shares (uncommitted)
hostname:configuration preferences> commit
hostname:configuration preferences> get login_screen
                 login_screen = shares
```

If you attempt to leave a context that contains uncommitted properties, you will be warned that leaving will abandon the set property values, and you will be prompted to confirm that you want to leave. For example:

```
hostname:configuration preferences> set login_screen=maintenance/hardware
                 login_screen = maintenance/hardware (uncommitted)
hostname:configuration preferences> done
You have uncommitted changes that will be discarded. Are you sure? (Y/N)
```

If a property in a context is set from a different context—that is, if the `set` command has been appended to a command that changes context—the commit is *implied*, and happens before control is returned to the originating context. For example:

```
hostname:> configuration preferences set login_screen=analytics/worksheets
                 login_screen = analytics/worksheets
hostname:>
```

Some properties take a list of values. For these properties, the list elements should be separated by a comma. For example, the NTP `servers` property may be set to a list of NTP servers:

```
hostname:configuration services ntp> set servers=0.pool.ntp.org,1.pool.ntp.org
                      servers = 0.pool.ntp.org,1.pool.ntp.org (uncommitted)
hostname:configuration services ntp> commit
```

If a property value contains a comma, an equals sign, a quote or a space, the entire value must be double quoted. For example, the `sharenfs` shares property for the default project may be set to read-only, but provide read/write access to host `kiowa`. For more information, see Shares and Projects.

```
hostname:> shares select default
hostname:shares default> set sharenfs="ro,rw=kiowa"
                     sharenfs = ro,rw=kiowa (uncommitted)
hostname:shares default> commit
```

Some properties are immutable; you can get their values, but you cannot set them. Attempts to set an immutable property results in an error. For example, attempting to set the immutable `space_available` property of the default project. For more information, see Shares and Projects.

```
hostname:> shares select default
hostname:shares default> get space_available
              space_available = 1.15T
hostname:shares default> set space_available=100P
error: cannot set immutable property "space_available"
```

Some other properties are only immutable in certain conditions. For these properties, the `set` command is not valid. For example, if the user named `pmd` is a network user, the `fullname` property will be immutable:

```
hostname:> configuration users select pmd set fullname="Rembrandt Q. Einstein"
error: cannot set immutable property "fullname"
```

Related Topics

- Browser User Interface (BUI)
- Command Line Interface (CLI)

Working with CLI Scripting

The CLI is designed to provide a powerful scripting environment for performing repetitive tasks.

You can use batch commands or scripting commands (or some combination), but in any case the automated infrastructure requires automated access to Oracle ZFS Storage Appliance. This must be done by user configuration, user authorizations, and setting SSH public keys using the CLI.

For information about configuring users, see the following:

- Configuring Users
- User Authorizations
- Setting SSH Public Keys (CLI)

To use CLI scripting, use the following sections:

- Using Batch Commands
- Understanding the CLI Scripting Commands
- Accessing the CLI Script Environment
- Understanding the Built-in CLI Functions
- Using the Children Function
- Using the Choices Function
- Using the Custom Alert Functions
- Using the Get Function
- Using the List Function
- Using the Prop Function
- Using the Run Function
- Using the Functions for Generating Output
- Understanding CLI Scripting Errors

Using Batch Commands

The simplest scripting mechanism is to batch appliance shell commands. For example, to automatically take a snapshot called `newsnap` in the project `myproj` and the filesystem `myfs`, put the following commands in a file:

```
shares
select myproj
select myfs
snapshots snapshot newsnap
```

Then `ssh` onto the appliance, redirecting standard input to be the file:

```
$ ssh root@hostname < myfile.txt
```

In many shells, you can abbreviate this by using a "here file", where input up to a token is sent to standard input. Following is the above example in terms of a here file:

```
$ '''ssh root@hostname << EOF
        shares
        select myproj
        select myfs
        snapshots snapshot newsnap
        EOF'''
```

This mechanism is sufficient for the simplest kind of automation, and may be sufficient if wrapped in programmatic logic in a higher-level shell scripting language on a client, but it generally leaves much to be desired.

Understanding the CLI Scripting Commands

While batch commands are sufficient for the simplest of operations, it can be tedious to wrap in programmatic logic. For example, if you want to get information on the space usage for every share, you must have many different invocations of the CLI, wrapped in a higher level language on the client that parsed the output of specific commands. This results in slow, brittle automation infrastructure. To allow for faster and most robust automation, the appliance has a rich *scripting environment* based on ECMAScript 3. An ECMAScript tutorial is beyond the scope of this document, but it is a dynamically typed language with a C-like syntax that allows for:

- Conditional code flow (`if`/`else`)
- Iterative code flow (`while`, `for`, and so on)
- Structural and array data manipulation via first-class Object and Array types
- Perl-like regular expressions and string manipulation (`split()`, `join()`, and so on)
- Exceptions
- Sophisticated functional language features, like closures

Accessing the CLI Script Environment

Use the following procedure to access the CLI script environment.

1. In the CLI, enter the script environment using the `script` command:

   ```
   hostname:> script
   ("." to run)>
   ```

2. At the script environment prompt, you can input your script, finally entering a period character (.) alone on a line to execute the script:

   ```
   hostname:> script
   ("." to run)> for (i = 10; i > 0; i--)
   ("." to run)>    printf("%d... ", i);
   ("." to run)> printf("Blastoff!\n");
   ("." to run)> .
   10... 9... 8... 7... 6... 5... 4... 3... 2... 1... Blastoff!
   ```

3. If your script is a single line, you can simply provide it as an argument to the script command, making for an easy way to explore scripting:

   ```
   hostname:> script print("It is now " + new Date())
   It is now Tue Oct 11 2022 05:33:01 GMT+0000 (UTC)
   ```

Understanding the Built-in CLI Functions

The following built-in functions enable your scripts to interact with the system.

Table 1-9 Built-in Functions to Support System Interactions

Function	Description
`children`	Returns an array of static children. See Using the Children Function.
`choices`	Returns an array of the valid property values for any property for which the set of values is known and enumerable. See Using the Choices Function.

`createalert,`
`postalert`
	Creates or posts a custom alert. See Using the Custom Alert Functions.

`get`	Gets the value of the specified property. Note that this function returns the value in native form. For example, dates are returned as Date objects. See Using the Get Function.
`list`	Returns an array of tokens corresponding to the dynamic children of the current context. See Using the List Function.
`prop`	Returns the value of the specified property in the current node context, or sets the specified property to a specified value. Both input and output values are in scriptable format. The values returned by `prop` might differ from values returned by `get`. See Using the Prop Function.
`props`	Returns an array of the property names for the current node.
`run`	Runs the specified command in the shell, returning any output as a string. Note that if the output contains multiple lines, the returned string will contain embedded newlines. See Using the Run Function.
`set`	Takes two string arguments, setting the specified property to the specified value.

Using the Children Function

Even in a context with static children, it can be useful to iterate over those children programmatically. This can be done by using the `children` function, which returns an array of static children.

1. For example, here is a script that iterates over every service, printing out the status of the service:

```
configuration services
script
       var svcs = children();
       for (var i = 0; i < svcs.length; ++i) {
                run(svcs[i]);
                try {
                        printf("%-10s %s\n", svcs[i], get('<status>'));
                } catch (err) { }
                run("done");
       }
```

2. Here is the output of running the script, assuming it was saved to a file named `svcinfo.aksh`:

```
$ ssh root@hostname < svcinfo.aksh
Password:
cifs       disabled
dns        online
ftp        disabled
http       disabled
identity   online
idmap      online
```

```
ipmp       online
iscsi      online
ldap       disabled
ndmp       online
nfs        online
nis        online
ntp        online
scrk       online
sftp       disabled
smtp       online
snmp       disabled
ssh        online
tags       online
vscan      disabled
```

Using the Choices Function

The `choices` function returns an array of the valid property values for any property for which the set of values is known and enumerable. For example, the following script retrieves the list of all pools on the shares node using the `choices` function and then iterates all pools to list projects and shares along with the available space.

1. For example, the following script retrieves the list of all pools on the shares node using the `choices` function, and then iterates all pools to list projects and shares along with the available space.

```
fmt = '%-40s %-15s %-15s\n';
printf(fmt, 'SHARE', 'USED', 'AVAILABLE');
run('cd /');
run('shares');
pools = choices('pool');
for (p = 0; p < pools.length; p++) {
        set('pool', pools[p]);
        projects = list();
        for (i = 0; i < projects.length; i++) {
                run('select ' + projects[i]);
                shares = list();
                for (j = 0; j < shares.length; j++) {
                        run('select ' + shares[j]);
                        share = pools[p] + ':' + projects[i] + '/' + shares[j];
                        printf(fmt, share, get('space_data'),
                            get('space_available'));
                        run('cd ..');
                }
                run('cd ..');
        }
}
```

2. Here is the output of running the script:

```
SHARE                                    USED            AVAILABLE
pond:projectA/fs1                        31744           566196178944
pond:projectA/fs2                        31744           566196178944
pond:projectB/lun1                       21474836480     587670999040
puddle:deptA/share1                      238475          467539219283
puddle:deptB/share1                      129564          467539219283
puddle:deptB/share2                      19283747        467539219283
```

Using the Custom Alert Functions

The `createalert` function creates a custom alert, and the `postalert` function posts the alert. The `createalert` function can be called from a script or from a workflow. The `postalert` function must be called from a workflow. See Creating and Posting Custom Alerts from Within a Workflow.

1. The following function call creates a minimalist custom alert.

   ```
   createalert([{'handler': 'syslog', 'args': {}}], {description: 'Writes to syslog'});
   ```

2. The following script specifies additional parameters and captures the return value.

   The return value is required to post the alert. Although the `postalert` call is shown in this script, `postalert` needs to be called from within a workflow in response to an event that occurs in the workflow.

   ```
   script
   ("." to run)> var actions = [{
   ("." to run)>     handler: 'email',
   ("." to run)>     args: {
   ("." to run)>         address: 'admin@example.com',
   ("." to run)>         subject: 'Custom Alert Response'
   ("." to run)>     }
   ("." to run)> }];
   ("." to run)> var params = {
   ("." to run)>     severity: 'Minor',
   ("." to run)>     description: 'Custom alert description',
   ("." to run)>     response: 'What the alert action does',
   ("." to run)>     impact: 'What happened to the appliance',
   ("." to run)>     recommended_action: 'What the administrator should do'
   ("." to run)> };
   ("." to run)> var cuuid = createalert(actions, params);
   ("." to run)> print(cuuid);
   ("." to run)> var puuid = postalert(cuuid);
   ("." to run)> .
   54c24732-b9c5-4b57-9aee-aeaf195afdae
   ```

Using the Get Function

The `run` function is sufficiently powerful that it may be tempting to rely exclusively on parsing output to get information about the system, but this has the decided disadvantage that it leaves scripts parsing human-readable output that may or may not change in the future. To more robustly gather information about the system, use the built-in `get` function. In the case of the `boot_time` property, this will return not the string but rather the ECMAScript `Date` object, allowing the property value to be manipulated programmatically. For more reliable scriptable values, see Using the Prop Function.

1. For example, you might want to use the `boot_time` property in conjunction with the current time to determine the time since boot:

   ```
   script
          run('configuration version');
          now = new Date();
          uptime = (now.valueOf() - get('boot_time').valueOf()) / 1000;
          printf('up %d day%s, %d hour%s, %d minute%s, %d second%s\n',
              d = uptime / 86400, d < 1 || d >= 2 ? 's' : '',
              h = (uptime / 3600) % 24, h < 1 || h >= 2 ? 's': '',
              m = (uptime / 60) % 60, m < 1 || m >= 2 ? 's': '',
              s = uptime % 60, s < 1 || s >= 2 ? 's': '');
   ```

2. Assuming the above is saved as `uptime.aksh`, you could run it this way:

   ```
   $ ssh root@hostname < uptime.aksh
   Pseudo-terminal will not be allocated because stdin is not a terminal.
   Password:
   up 2 days, 10 hours, 47 minutes, 48 seconds
   ```

   The message about pseudo-terminal allocation is due to the SSH client; the issue that this message refers to can be managed by specifying the `T` option to SSH.

Using the List Function

In a context with dynamic children, it can be very useful to iterate over those children programmatically. This can be done by using the `list` function, which returns an array of dynamic children.

1. The following example script iterates over every share in every project, printing out the amount of space consumed and space available:

   ```
   script
          run('shares');
          projects = list();

          for (i = 0; i < projects.length; i++) {
                  run('select ' + projects[i]);
                  shares = list();

                  for (j = 0; j < shares.length; j++) {
                          run('select ' + shares[j]);
                          printf("%s/%s %1.64g %1.64g\n", projects[i], shares[j],
                              get('space_data'), get('space_available'));
                          run('cd ..');
                  }

                  run('cd ..');
          }
   ```

2. Here is the output of running the script, assuming it was saved to a file named `space.aksh`:

   ```
   $ ssh root@hostname < space.aksh
   Password:
   admin/accounts 18432 266617007104
   admin/exports 18432 266617007104
   admin/primary 18432 266617007104
   admin/traffic 18432 266617007104
   admin/workflow 18432 266617007104
   aleventhal/hw_eng 18432 266617007104
   bcantrill/analytx 1073964032 266617007104
   bgregg/dashbd 18432 266617007104
   bgregg/filesys01 26112 107374156288
   bpijewski/access_ctrl 18432 266617007104
   ...
   ```

3. If you would prefer a "pretty printed" (although more difficult to handle programmatically) variant of this, you could directly parse the output of the `get` command:

   ```
   script
          run('shares');
          projects = list();

          printf('%-40s %-10s %-10s\n', 'SHARE', 'USED', 'AVAILABLE');
   ```

```
for (i = 0; i < projects.length; i++) {
        run('select ' + projects[i]);
        shares = list();

        for (j = 0; j < shares.length; j++) {
                run('select ' + shares[j]);

                share = projects[i] + '/' + shares[j];
                used = run('get space_data').split(/\s+/)[3];
                avail = run('get space_available').split(/\s+/)[3];

                printf('%-40s %-10s %-10s\n', share, used, avail);
                run('cd ..');
        }

        run('cd ..');
}
```

4. Here is the output of running this new script, assuming it was named `prettyspace.aksh`:

   ```
   $ ssh root@hostname < prettyspace.aksh
   Password:
   SHARE                                    USED       AVAILABLE
   admin/accounts                           18K        248G
   admin/exports                            18K        248G
   admin/primary                            18K        248G
   admin/traffic                            18K        248G
   admin/workflow                           18K        248G
   aleventhal/hw_eng                        18K        248G
   bcantrill/analytx                        1.00G      248G
   bgregg/dashbd                            18K        248G
   bgregg/filesys01                         25.5K      100G
   bpijewski/access_ctrl                    18K        248G
   ...
   ```

5. The `list` function supports optional arguments `depth` and `filter`.

   The format is: `list ([depth, [filter]])`. The argument `depth` can be defined by a number. The greater number of `depth`, the more details will be returned. The argument `filter` is formatted as `{<prop1>:<val1>, <prop2>:<val2> ...}`. If `filter` is specified, `depth` must also be specified.

   Usage and input behavior:

   - `list()` - Returns only node names.
   - `list(0)` - Return properties of node and only children names.
   - `list(0, {kiosk_mode: true})` - Return a filtered list for `kiosk_mode` is `true` with names of children.
   - `list(1)` - Return properties of node, names and properties of children, only names of grandchildren.
   - `list(1, {kiosk_mode: true})` - Return a filtered list for `kiosk_mode` is `true` with details up to `depth=1`.
   - `list(2)` - Return properties of node, names and properties of children and `list(0)` output of grandchildren.
   - `list(2, {fullname:'Super*', kiosk_mode: true})` - Return a filtered list for `fullname` containing `Super` and `kiosk_mode` is `true` with details up to `depth=2`.

6. This is an example output for a list with `depth=2`:

   The label `name` shows the name of the list item (that is, a node). The label `properties` shows the properties of the list item. The label `children` shows static children of the list item. The label `list` shows dynamic children of the list item.

```
script
        ("." to run)> dump(list(2));
        ("." to run)> .

        [{
            name: 'restuser',
            properties: {
                kiosk_screen: 'status/dashboard',
                kiosk_mode: false,
                roles: ['basic'],
                require_annotation: false,
                initial_password: 'DummyPassword',
                fullname: 'REST User',
                logname: 'restuser'
            },
            children: [{
                name: 'preferences',
                properties: {
                    advanced_analytics: false,
                    session_timeout: 15,
                    login_screen: 'status/dashboard',
                    locale: 'C'
                }
            }, {
                name: 'exceptions',
                list: [{
                    name: 'auth-000',
                    properties: {
                         allow_configure: false,
                         scope: 'alert'
                    }
                }, {
                    name: 'auth-001',
                    properties: {
                         allow_workgroup: false,
                         allow_domain: false,
                         name: '*',
                         scope: 'ad'
                    }
                }]
            }]
        }]
```

Using the Prop Function

In most cases, the `get` and `prop` functions return the same value. The difference between these two functions is that the `get` function does not always return a scriptable value, while the `prop` function always returns a scriptable value. A scriptable value has the same stable form for each type of property.

1. Return the scriptable value of a property in the current node context.

   For most properties, the `get` and `prop` functions return identical values, as shown in the following example for the `version` property:

```
> ls
version = 2019.02.28,1-0
> script get('version')
'2019.02.28,1-0'
> script prop('version')
'2019.02.28,1-0'
```

In the following example for the `date` property, note the difference between the value returned by the `get` function and the scriptable value returned by the `prop` function:

```
> ls
date = 2019-2-28 10:43:11
> script get('date')
Thu Feb 28 2019 10:43:11 GMT+0000 (UTC) (Date object)
> script prop('date')
'2019-02-28T10:43:11Z'
```

2. Set the value of a property in the current node context to a scriptable value.

   Include a value in the `prop` function call to set the value of the named property. An error message is output if the given value is not in the specified scriptable form for that property. Otherwise, the set form of the `prop` function does not return any value.

   ```
   > ls
   date = 2019-2-28 10:43:11
   > script prop('date', '2019-03-09T12:34:56Z')
   > ls
   date = 2019-3-09 12:34:56
   > script prop('date')
   '2019-03-09T12:34:56Z'
   ```

3. Set the value of a boolean property.

   Because the `prop` function can be used either to return a property value or to set a property value, the `prop` function does not have the ability that the `set` function has to set a boolean property to `true` without specifying the value.

   The following example shows how to use the short version of the `set` function to set a boolean property to the value `true`:

   ```
   set('booleanproperty')
   ```

   The short form of the `prop` function returns the value of the boolean property, either `true` or `false`:

   ```
   script prop('booleanproperty')
   'false'
   ```

   To use the `prop` function to set the value of a boolean property, you must provide the value:

   ```
   script prop('booleanproperty', true)
   script prop('booleanproperty')
   'true'
   ```

4. Set the value of a `List` property.

   To set a scriptable value for a property with the `List` modifier, you must specify an array of values:

   ```
   > ls
   stringlist = a,string,list
   > script prop('stringlist', ['a', 'b', 'c', 'd'])
   > script prop('stringlist')
   ```

```
['a', 'b', 'c', 'd']
> ls
stringlist = a,b,c,d
```

To specify a single value for a `List` property, specify the single value as an array of size one:

```
> ls
stringlist = a,string,list
> script prop('stringlist', ['a'])
> script prop('stringlist')
['a']
> ls
stringlist = a
```

To specify the empty value for a property with the `Empty` or `List` modifier, specify an empty array:

```
> ls
emptystringlist = a,string,list
> script prop('emptystringlist', [])
> script prop('emptystringlist')
[]
> ls
emptystringlist =
```

Using the Run Function

1. The simplest way for scripts to interact with the larger system is to use the `run` function. The `run` function takes a command to run, and returns the output of that command as a string. For example:

```
hostname:> configuration version script dump(run('get boot_time'))
'                      boot_time = 2022-10-12 07:02:17\n'
```

2. The built-in dump function dumps the argument out, without expanding any embedded newlines. ECMAScript's string handling facilities can be used to take apart output. For example, splitting the previous example based on whitespace (negative space):

```
hostname:> configuration version script dump(run('get
boot_time').split(/\s+/))
['', 'boot_time', '=', '2022-10-12', '07:02:17', '']
```

Using the Functions for Generating Output

Reporting state on the system requires generating output. Scripts have several built-in functions made available to them to generate output.

Table 1-10 Built-in Functions for Generating Output

Function	Description
`dump`	Dumps the specified argument to the terminal, without expanding embedded newlines. Objects will be displayed in a JSON-like format. Useful for debugging.
`print`	Prints the specified object as a string, followed by a newline. If the object does not have a `toString` method, it will be printed opaquely.
`printf`	Like C's `printf(3C)`, prints the specified arguments according to the specified formatting string.

Understanding CLI Scripting Errors

When an error is generated, an exception is thrown. The exception is generally an object that contains the following members:

- `code` - a numeric code associated with the error
- `message` - a human-readable message associated with the error

Exceptions can be caught and handled, or they may be thrown out of the script environment. If a script environment has an uncaught exception, the CLI will display the details. For example:

```
hostname:> script run('not a cmd')
error: uncaught error exception (code EAKSH_BADCMD) in script: invalid command
       "not a cmd" (encountered while attempting to run command "not a cmd")
```

You could see more details about the exception by catching it and dumping it out:

```
hostname:> script try { run('not a cmd') } catch (err) { dump(err); }
{
   toString: <function>,
   code: 10004,
   message: 'invalid command "not a cmd" (encountered while attempting to
                       run command "not a cmd")'
}
```

This also allows you to have rich error handling; for example:

```
#!/usr/bin/ksh -p

ssh -T root@hostname <<EOF
script
       try {
               run('shares select default select $1');
       } catch (err) {
               if (err.code == EAKSH_ENTITY_BADSELECT) {
                       printf('error: "$1" is not a share in the ' +
                           'default project\n');
                       exit(1);
               }

               throw (err);
       }

       printf('"default/$1": compression is %s\n', get('compression'));
       exit(0);
EOF
```

If this script is named `share.ksh` and run with an invalid share name, a rich error message will be generated:

```
$ ksh ./share.ksh bogus
error: "bogus" is not a share in the default project
```

2 Configuring the Appliance

To configure Oracle ZFS Storage Appliance, use the following sections:

- Initial Appliance Configuration
- Appliance Cluster Configuration
- Network Configuration
- Configuring Storage
- Configuring Cloud Backup
- Understanding the Appliance Status
- Configuring Storage Area Network (SAN)
- Configuring Users
- Setting Appliance Preferences
- Configuring Alerts
- Configuring Certificates
- Configuring SSL/TLS Versions and Ciphers
- Configuring Password Complexity

Initial Appliance Configuration

If you are setting up a new Oracle ZFS Storage Appliance, follow the initial configuration steps in Configuring the Appliance for the First Time in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*.

You can repeat initial configuration at a later time by clicking the **INITIAL SETUP** button on the **Maintenance: System** screen, or by entering the `maintenance system setup` context in the CLI.

Related Topics

- Appliance Cluster Configuration
- Network Configuration
- Configuring Storage

Appliance Cluster Configuration

Oracle ZFS Storage Appliance supports cooperative clustering of two controllers. Clustering controllers can be part of an integrated approach to enhancing availability that might also include client-side load balancing, proper site planning, proactive and reactive maintenance and repair, and the single-appliance hardware redundancy that is built into all appliances.

> **Note:**
>
> If you are configuring clustering for two new controllers, follow the BUI procedure for initial configuration of the new controllers, as described in Configuring the Appliance for the First Time in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*.

For procedures related to appliance clustering, see:

- Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*
- Cluster Configuration BUI View
- Checking Cluster Link Status (CLI)
- Upgrading a Standalone Appliance to a Clustered Configuration (BUI)
- Shutting Down a Clustered Configuration - BUI, CLI
- Unconfiguring a Cluster Node

For a better understanding of appliance clustering, see:

- Cluster Terminology
- Understanding Clustering
- Clustered Controller States
- Cluster Interconnect I/O
- Cluster Resource Management
- Cluster Takeover and Failback
- Configuration Changes in a Clustered Environment
- Clustering Considerations for Storage
- Clustering Considerations for Networking
- Private Local IP Interfaces
- Clustering Considerations for InfiniBand
- Preventing Split-Brain Conditions
- Estimating and Reducing Takeover Impact

Cluster Configuration BUI View

The **Configuration: Cluster** view provides a graphical overview of the status of the cluster card, the cluster controller states, and the cluster resources.

The following figure shows cluster connections between two Oracle ZFS Storage ZS11-2 controllers:

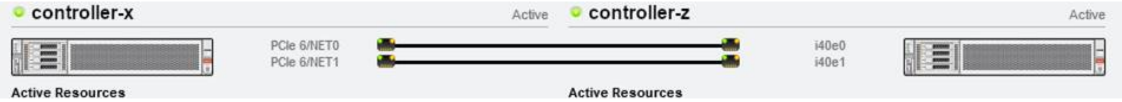

The following figure shows cluster connections between two Oracle ZFS Storage ZS9-2 controllers:

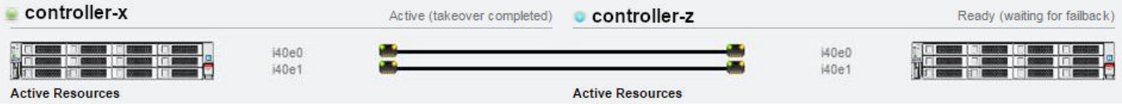

The following figure shows cluster connections between two Oracle ZFS Storage ZS7-2, ZS5-x, ZS4-4, ZS3-4, or Sun ZFS Storage 7x20 controllers:

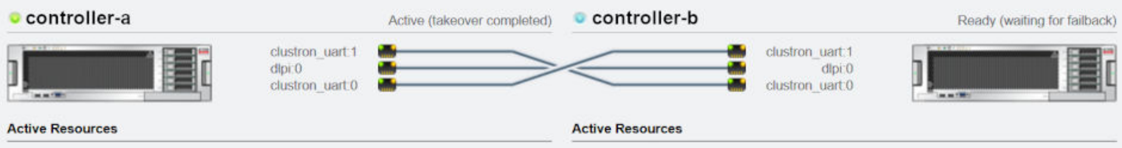

The following figure shows cluster connections between two Oracle ZFS Storage ZS3-2 controllers:

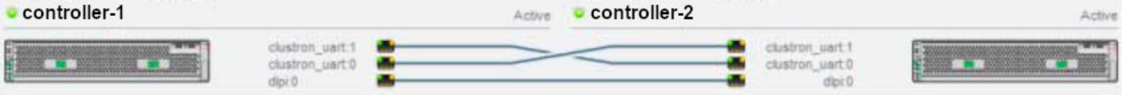

The interface contains the following objects:

- A thumbnail picture of each controller. The controller that the user is currently logged into is on the left side of the view. Above each controller thumbnail, the controller name is on the left, and the current cluster state of the controller (such as **Active** or **Ready**) is on the right. For descriptions of controller states, see Clustered Controller States.
- A line that represents each cluster card connection. These lines update dynamically with the hardware. A solid line indicates that the link is connected and active. A dashed line indicates that the link is broken in one of the following ways:
  - The other controller is restarting/rebooting.
  - The link is not cabled correctly or the cluster cables are not secure in their connectors. For cluster cabling instructions, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

  Ensure that all links are connected and active before you perform initial cluster setup. See Upgrading a Standalone Appliance to a Clustered Configuration (BUI).
- Below each controller thumbnail is a list of the **PRIVATE** and **SINGLETON** resources that are currently assigned to that controller, and some attributes of those resources such as IP address or size.
- For each resource, the owner of the resource is shown. The owner of the resource is the controller that will provide the resource when both controllers are in the **CLUSTERED** state. To change the owner, click the current owner name, select the peer controller, and then click the **APPLY** button at the top right of the view.
- To the right of the owner is a restart icon ⟳ that enables you to attempt to repair the resource.
- For each resource, a lock icon 🔒 indicates whether the resource is **PRIVATE**. When the current controller is in either the **OWNER** or **CLUSTERED** state, a resource can be locked

to it (made **PRIVATE**) or unlocked (made a **SINGLETON**). To lock a resource to the controller or unlock the resource from the controller, click the lock icon, and then click the **APPLY** button at the top right of the view. Note that **PRIVATE** resources that belong to the remote peer are not displayed on either resource list.

The following table describes the buttons at the top of the **Configuration: Cluster** view.

Table 2-1 Cluster Interface Buttons

Button	Description
Setup	The setup operation is a step in initial cluster configuration. See Upgrading a Standalone Appliance to a Clustered Configuration (BUI).
Unconfig	The unconfig operation configures a cluster node to standalone operation. See Unconfiguring a Cluster Node.
Failback	The failback operation changes the cluster configuration from OWNER-STRIPPED (active-passive) to CLUSTERED-CLUSTERED (active-active). See Upgrading a Standalone Appliance to a Clustered Configuration (BUI) and Cluster Takeover and Failback.
Takeover	Takeover is performed automatically in certain situations. Takeover can be performed manually, which can be useful for testing purposes. See Cluster Takeover and Failback.
Revert	If resource modifications are pending (resource rows are highlighted in yellow), revert those changes and show the current cluster configuration.
Apply	If resource modifications are pending (resource rows are highlighted in yellow), commit those changes to the cluster.

Related Topics

- Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*
- Upgrading a Standalone Appliance to a Clustered Configuration (BUI)
- Clustered Controller States

Checking Cluster Link Status (CLI)

Use this procedure to determine whether the cluster links are connected and active.

Before you perform initial cluster setup (see Upgrading a Standalone Appliance to a Clustered Configuration (BUI)), ensure that all cluster links are connected and active. A broken connection can mean that the other cluster node is restarting/rebooting, or it can mean that the link is not cabled correctly or the cluster cables are not secure in their connectors. For cluster cabling instructions, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

1. Go to `configuration cluster`.
2. Enter the `links` command.

   In the following examples, all links are connected and active (`AKCIOS_ACTIVE` state). A link that is not up will show a different state, such as `AKCIOS_TIMEDOUT`.

   The following example shows cluster connections between two Oracle ZFS Storage ZS11-2 or ZS9-2 controllers.

   ```
   hostname:configuration cluster> links
   ```

```
    PCIe 6/NET0 = AKCIOS_ACTIVE
    PCIe 6/NET1 = AKCIOS_ACTIVE
```

The following example shows cluster connections between two Oracle ZFS Storage ZS7-2 or ZS5-2 controllers.

```
hostname:configuration cluster> links

    clustron3_ng3:0/clustron_uart:0 = AKCIOS_ACTIVE
    clustron3_ng3:0/clustron_uart:1 = AKCIOS_ACTIVE
    clustron3_ng3:0/dlpi:0 = AKCIOS_ACTIVE
```

Other controllers show similar `links` output. The only difference is in the portion that precedes `/clustron` or `/dlpi`. For example, `clustron3_ng3:0` in the preceding example is `clustron3:0` for Oracle ZFS Storage ZS5-4 controllers, `clustron2:0` for Oracle ZFS Storage ZS4-4 controllers, and `clustron2_embedded:0` for Oracle ZFS Storage ZS3-2 controllers.

Related Topics

- For instructions about how to see the state of each controller, see Shutting Down a Clustered Configuration (CLI).
- For descriptions of controller states, see Clustered Controller States.

Upgrading a Standalone Appliance to a Clustered Configuration (BUI)

Use this procedure to configure a standalone appliance to be a node in a clustered configuration.

> **Note:**
>
> The BUI is strongly recommended as the interface to use to configure clustered controllers.

Before You Begin

Ensure the following:

- Both controllers that will be in the new cluster are the same model. The Sun ZFS Storage 7420 with 2GHz or 2.40GHz CPUs can be clustered with the Sun ZFS Storage 7420 with 1.86GHz or 2.00GHz CPUs.
- The second controller is a new controller or a controller that has been reset to factory settings. See Performing a Factory Reset in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.
- The standalone appliance is powered on. You do not need to power down the standalone appliance during this procedure.

1. Connect the cluster cables between the standalone appliance and second controller.

   For cluster cabling instructions, see Performing a Factory Reset in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

2. On the second controller, connect the power cables into power supply 0 and power supply 1. Then connect each cable to the external power source.

   The second controller powers on automatically.

3. Connect the second controller to the disk shelves.

   See the documentation that came with your appliance or refer to Getting Started with Cabling in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

4. On the standalone controller, from the **Configuration** menu, select **Cluster**.

5. Confirm that the communication links between the two controllers are connected and active.

   All cluster connections should be solid lines as shown in Cluster Configuration BUI View, or should be in state `AKCIOS_ACTIVE` as shown in Checking Cluster Link Status (CLI). A broken connection can mean that the other controller is restarting/rebooting, or it can mean that the link is not cabled correctly or the cluster cables are not secure in their connectors. For cluster cabling instructions, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

6. Click **SETUP**.

7. Enter the host name for the second controller and the same root password that is set on the first controller.

   > **Note:**
   >
   > Initial cluster configuration setup can take several minutes to complete.

8. On the standalone controller, from the **Configuration** menu, select **Cluster** and lock the management interface.

   Click the lock icon 🔒 for the management interface to lock that interface.

   Locking the management interface to the controller will prevent a transfer of resources when a failback occurs.

9. From the standalone controller, configure the management interface for the second controller.

   a. From the **Configuration** menu, select **Network**, and click the add icon ⊕ next to **Interfaces**.

   b. Enter a name for the management interface, and check the boxes for **Enable Interface and Allow Administration**.

   c. Select an IP address, and click **APPLY**.

10. From the **Configuration** menu, select **Cluster**, and click **FAILBACK** to bring the cluster to **Active:Active** mode.

    The two controllers are now configured as clustered peers.

11. On the second controller, from the **Configuration** menu, select **Cluster**, and lock the management interface.

    Click the lock icon 🔒 for the management interface to lock that interface.

Related Topics

- Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*
- Cluster Configuration BUI View
- Checking Cluster Link Status (CLI)
- Cluster Takeover and Failback

Shutting Down a Clustered Configuration (BUI)

Use this procedure to shut down a clustered configuration. A cluster that has been shut down can be restored to the same cluster configuration as before the shutdown by powering on both controllers.

1. From one of the peer controllers: From the **Configuration** menu, select **Cluster**.
2. Determine the cluster state of each controller.

   In the following figure, the active controller is `controller-a`, and the standby controller is `controller-b`.

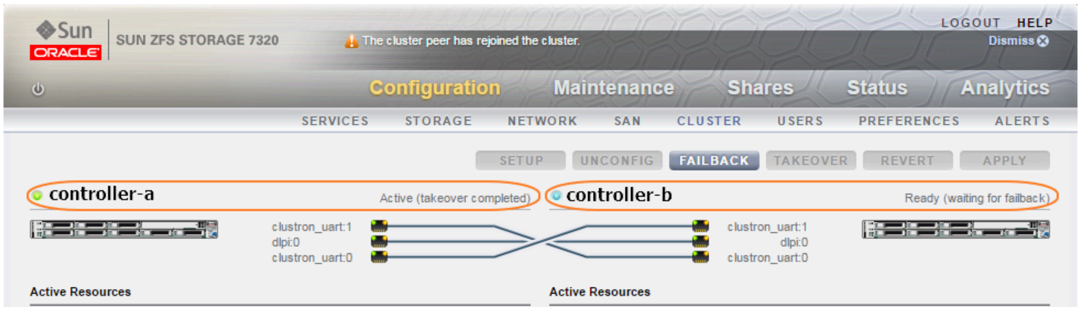

   The following table describes the valid pairs of controller states (active and standby controllers) that you could see.

controller-a	controller-b	Condition
Active	Active	Both controllers are running in a normal clustered condition.
Active (takeover completed)	● Ready (waiting for failback)	`controller-a` owns all of the resources and is the active controller. `controller-b` is in standby mode and has no resources. To limit the number of times a pool is moved, shut down the standby controller first.
Active (takeover completed)	◎ Rejoining cluster ...	`controller-b` is rebooting and `controller-a` has all resources.
Active (takeover completed)	Unknown (disconnected or restarting)	`controller-b` is powered off or rebooting, or all of its cluster interconnect links are down.

3. Log in to the BUI of `controller-b`, and click the power icon ⏻ on the left side, under the masthead.

> **Note:**
>
> To limit the number of times a pool is moved, shut down the standby controller first.

4. From the BUI of `controller-a's` **Configuration** menu, select **Cluster** to confirm that `controller-b` is powered off, with the cluster state: **Unknown (disconnected or restarting)**.

5. From the BUI of `controller-a`, click the power icon ⏻ on the left side, under the masthead.

6. Optional: Use Oracle ILOM to confirm that both controllers are powered off.

   Enter the following command in the Oracle ILOM:

   ```
   ->show /SYS power_state
   ```

   For information about accessing Oracle ILOM, see Logging in to Oracle ILOM Remotely Using a Command Line Interface in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

7. Power off the disk shelves.

   a. Place the power supply on/off switches to the "O" off position.

   b. Disconnect the power cords from the external power source for the cabinet.

   > **Note:**
   >
   > All power cords must be disconnected to completely remove power from the disk shelf.

   For more information, see Powering Off a Disk Shelf in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*.

Related Topics

- Understanding Clustering
- Cluster Resource Management
- Cluster Takeover and Failback
- Configuration Changes in a Clustered Environment

Shutting Down a Clustered Configuration (CLI)

Use this procedure to shut down a clustered configuration. A cluster that has been shut down can be restored to the same cluster configuration as before the shutdown by powering on both controllers.

1. Go to `configuration cluster`.

2. Determine the cluster state of each controller.

   In the following example, `controller-a` is the owner and in the `active` state. Its peer, `controller-b`, is the standby controller and in the `stripped` state.

```
controller-a:> configuration cluster
controller-a:configuration cluster> get
                         state = AKCS_OWNER
                   description = Active (takeover completed)
                      peer_asn = 365ed33c-3b9d-c533-9349-8014e9da0408
                 peer_hostname = controller-b
                    peer_state = AKCS_STRIPPED
              peer_description = Ready (waiting for failback)
```

The following table describes the pairs of controller states (`state` and `peer_state`) that you could see.

controller-a	controller-b	Condition
`AKCS_CLUSTERED`	`AKCS_CLUSTERED`	Both controllers are running in a normal clustered condition.
`AKCS_OWNER`	`AKCS_STRIPPED`	`controller-a` owns all of the resources and is the active controller. `controller-b` is in standby mode and has no resources. To limit the number of times a pool is moved, shut down the `STRIPPED` controller first.
`AKCS_OWNER`	`REBOOTING`	`controller-a` has all resources. `controller-b` is rebooting.
`AKCS_OWNER`	`UNKNOWN`	`controller-a` has all resources. `controller-b` is powered off or rebooting, or all of its cluster interconnect links are down.

3. Shut down `controller-b`.

```
controller-b:configuration cluster> cd /
controller-b:> maintenance system poweroff
This will turn off power to the appliance. Are you sure? (Y/N) Y
```

> **Note:**
>
> If both controllers have a status of `AKCS_CLUSTERED`, a takeover of the surviving controller begins automatically.

4. From `controller-a`, verify that `controller-b` has been powered off and is in state `OWNER/unknown`.

```
controller-a:configuration cluster> get
                         state = AKCS_OWNER
                   description = Active (takeover completed)
                      peer_asn = 365ed33c-3b9d-c533-9349-8014e9da0408
                 peer_hostname = controller-b
                    peer_state = OWNER/unknown
              peer_description =
```

5. Shut down `controller-a`.

```
controller-a:configuration cluster> cd /
controller-a:> maintenance system poweroff
This will turn off power to the appliance. Are you sure? (Y/N) Y
```

6. Optional: Use Oracle ILOM to confirm that both controllers are powered off.

Enter the following command in the Oracle ILOM CLI:

```
->show /SYS power_state
```

For information about accessing Oracle ILOM, see Logging in to Oracle ILOM Remotely Using a Command Line Interface in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

7. Power off the disk shelves.
   a. Place the power supply on/off switches to the "O" off position.
   b. Disconnect the power cords from the external power source for the cabinet.

   > **Note:**
   >
   > All power cords must be disconnected to completely remove power from the disk shelf.

   For more information, see Powering Off a Disk Shelf in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*.

Related Topics

- Understanding Clustering
- Cluster Resource Management
- Cluster Takeover and Failback
- Configuration Changes in a Clustered Environment

Unconfiguring a Cluster Node

This procedure describes how to unconfigure a cluster node to standalone operation. A cluster node that has been unconfigured cannot be restored to the same cluster configuration because the unconfigure operation is destructive and results in data loss.

> **Caution:**
>
> In general, do not unconfigure a cluster node yourself. Unconfiguring a cluster node results in data loss. Contact Oracle Support.

The following are reasons that you might want to unconfigure clustering:

- You no longer want to use clustering. Instead, you want to configure two independent storage appliances.
- You are replacing a failed storage controller with:
  - New hardware
  - A storage controller with factory-fresh appliance software

> **Note:**
>
> This replacement should be performed by your service provider.

1. Back up the data on the cluster storage resources.
2. On the first controller (`controller-a`), perform a factory reset.

   See Factory Resetting Clustered Controllers in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.
3. When `controller-a` starts to reboot, immediately turn off its power.

   Verify the power off and controller state. See Clustered Controller States.
4. Completely disconnect `controller-a`, including all network, SAS, and clustering cables.
5. On the second controller (`controller-b`), from the **Configuration** BUI menu, select **Cluster**, and click **UNCONFIG**.

   `Controller-b` can still serve data.
6. Optional: Reconfigure the cluster.

   `Controller-b` is now a standalone appliance, and `controller-a` is factory reset and powered off. If you want to reconfigure these into a cluster, follow the instructions in Upgrading a Standalone Appliance to a Clustered Configuration (BUI).

Cluster Terminology

Cluster states and resource types are described in Understanding Clustering.

- **Export** - Make a resource inactive on a particular controller.
- **Failback** - Move from `AKCS_OWNER` state to `AKCS_CLUSTERED` state, in which all foreign resources (those assigned to the peer) are exported, then imported by the peer.
- **Import** - Make a resource active on a particular controller.
- **Peer** - The other controller in a cluster, sometimes also called the standby controller.
- **Rejoin** - Retrieve and resynchronize the resource map from the peer.
- **Resource** - A physical or virtual object present, and possibly active, on one or both controllers.
- **Takeover** - Move from `AKCS_CLUSTERED` or `AKCS_STRIPPED` state to `AKCS_OWNER` state, in which all resources are imported.

Understanding Clustering

A cluster is a metasystem composed of two appliance controllers and shared storage. The clustering subsystem consists of three main building blocks:

- **Cluster I/O subsystem** - The cluster I/O subsystem and the hardware device (cluster I/O ports) provide a transport for inter-controller communication within the cluster and are responsible for monitoring each peer's state. See Cluster Interconnect I/O.
- **Resource manager** - The transport is used by the resource manager, which allows data service providers and other management subsystems to interface with the clustering system. See Cluster Resource Management.

- **Cluster management user interfaces** - The cluster management user interfaces provide setup, resource allocation and assignment, monitoring, and takeover and failback operations.

A cluster improves availability and reduces business disruption in the following ways:

- If one of the controllers experiences certain hardware or software failures, the second controller provides service while repair or replacement of the first controller is performed.
- Rolling software upgrade can reduce the business disruption associated with migrating to newer software.

Each controller can be assigned storage, networking, and other resources from the set of resources that is available to the cluster. A cluster has one of the following topologies, depending on how resources are assigned to the controllers:

- **active-active** - Each controller in the cluster is assigned at least one storage pool and the network resources needed by clients to reach the data stored in that pool.
- **active-passive** - A single storage pool and the network resources needed by clients to reach the data stored in that pool is assigned to the controller that is designated as the *active* controller.

If a controller fails, the other controller (the *peer*) takes control of all known resources and provides the services associated with those resources.

Clustered Controller States

Clustered controller nodes are in one of the following states at any given time.

Table 2-2 Cluster States

BUI/CLI State	Description
● Clustering is not configured `AKCS_UNCONFIGURED`	Clustering has not yet been configured. The cluster is either being set up or the cluster setup task has never been completed. The peer controller typically is in state `UNKNOWN` (disconnected or restarting).
Unknown (disconnected or restarting) `UNKNOWN`	The controller is powered off or rebooting, all of its cluster interconnect links are down, or clustering has not yet been configured.
Active (takeover completed) `AKCS_OWNER`	Clustering is configured, and this controller has taken control of all shared resources in the cluster. A controller enters this state in the following cases: • Immediately after cluster setup is completed from this controller's user interface. • When this controller detects that its peer has failed, for example after a takeover. The controller remains in this state until an administrator manually executes a failback operation. See Cluster Takeover and Failback.

Table 2-2 (Cont.) Cluster States

BUI/CLI State	Description
● Ready (waiting for failback) `AKCS_STRIPPED`	Clustering is configured, and this controller does not control any shared resources. A controller enters this state in the following cases: • Immediately after cluster setup is completed from the user interface of this controller's peer controller. • Following a reboot, power disconnect, or other failure. The controller remains in this state until an administrator manually executes a failback operation. See Cluster Takeover and Failback.
Active `CLUSTERED`	Clustering is configured, and both nodes own shared resources according to their resource assignments. If each node owns a ZFS pool and is in the `CLUSTERED` state, then the two nodes form what is commonly called an *active-active cluster*.
◉ Rejoining cluster ... `REBOOTING`	The appliance has recently rebooted, or the appliance management software is restarting after an internal failure. Resource state is being resynchronized.

Cluster Interconnect I/O

All inter-controller communication consists of one or more messages transmitted over the redundant cluster I/O links provided by the controller cluster interface card. For more information about cluster interface cards and cluster cabling, see Controller Cluster I/O Ports in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x* and Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

- Oracle ZFS Storage ZS11-2 and ZS9-2 controllers employ Ethernet-based clustering using two Ethernet ports in the Oracle Quad Port 10GBASE-T Ethernet Adapter.
- Oracle ZFS Storage ZS7-2, ZS5-x, ZS4-4, ZS3-x, and Sun ZFS Storage 7x20 controllers employ serial-based clustering using two serial cluster links, and provide Ethernet connectivity via one link. The Ethernet link provides a higher-performance transport for non-heartbeat messages, such as rejoin synchronization, and provides a backup heartbeat.

Clustered controllers only communicate with each other over the secure private network established by the cluster interconnects, never over network interfaces intended for service or administration. Messages fall into two general categories: regular heartbeats used to detect the failure of a remote controller, and higher-level traffic associated with the resource manager and the cluster management subsystem.

Heartbeats are sent, and expected, on all links. Heartbeats are transmitted continuously at fixed intervals. Heartbeats are never acknowledged or retransmitted because all heartbeats are identical and contain no unique information. Other traffic is acknowledged, verified, and retransmitted as required to maintain a reliable transport for higher-level software.

For Oracle ZFS Storage ZS11-2 and ZS9-2 controllers, heartbeat messages are sent at 200ms intervals. Failure to receive any message after 1 second is considered to be link failure. For all other controllers, heartbeat messages are sent on all cluster I/O links at 50ms intervals. Failure to receive any message after 200ms (serial links) or 500ms (Ethernet links) is considered to be link failure. For all controllers, if all links have failed, the peer is assumed to have failed, and takeover arbitration will be performed.

If a panic occurs on an Oracle ZFS Storage ZS11-2 or ZS9-2 controller, the clustering system can detect that the peer has failed within 1200ms. No panic message is sent.

If a panic occurs on a, Oracle ZFS Storage ZS7-2, ZS5-x, ZS4-4, ZS3-x, or Sun ZFS Storage 7x20 controller, the panicking controller will transmit a single notification message over each serial link. The peer controller will immediately begin takeover, regardless of the state of any other links. Given these characteristics, the clustering subsystem normally can detect that the peer has failed within:

- 550ms, if the peer has stopped responding or lost power, or
- 30ms, if the peer has encountered a fatal software error that triggered an operating system panic.

All of the values described in this section are fixed. The appliance does not offer the ability to tune these parameters. These parameters are provided here for informational purposes only and may be changed without notice at any time.

> **Note:**
>
> To avoid data corruption after a physical re-location of a cluster, verify that all cluster cabling is installed correctly in the new location. For more information, see Preventing Split-Brain Conditions.

Related Topics

Clustered Controller States

Cluster Resource Management

The resource manager is responsible for ensuring that the correct set of network interfaces is plumbed up, the correct storage pools are active, and the numerous configuration parameters remain in sync between two clustered controllers. Most of this subsystem's activities are invisible to administrators. However, one important aspect is exposed. Resources are classified into several types that govern when and whether the resource is imported (made active). Note that the definition of active varies by resource class. For example, a network interface belongs to the net class and is active when the interface is brought up.

The three most important resource types are replica, singleton, and private:

- **Replica resources -** Replicas are simplest: They are never exposed to administrators and do not appear on the cluster configuration screen. Replicas always exist and are always active on both controllers. Typically, these resources simply act as containers for service properties that must be synchronized between the two controllers.
- **Singleton resources -** Like replicas, singleton resources provide synchronization of state. However, singletons are always active on exactly one controller. Administrators can choose the controller on which each singleton should normally be active. If that controller has failed, its peer will import the singleton. Singletons are the key to the availability characteristics of clustering. Singletons are the resources one typically imagines moving from a failed controller to its surviving peer. Singletons include network interfaces and storage pools. Because a network interface is a collection of IP addresses used by clients to find a known set of storage services, it is critical that each interface be assigned to the

same controller that the storage pool clients will expect to see when accessing that interface's addresses.

- **Private resources -** Private resources are known only to the controller to which they are assigned, and are never taken over upon failure. This is typically useful only for network interfaces. See the following discussion of specific use cases.

Another data type is the *symbiote*. A symbiote allows one resource to follow another as it is imported and exported. For example, a symbiote is used to represent disks and flash devices in the storage pool.

The following figure shows an example of a clustered configuration, including shared resources.

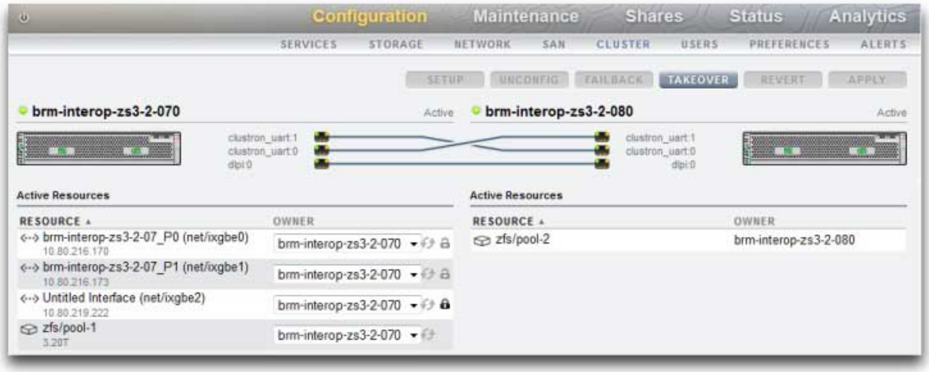

Active Resources

RESOURCE ▲	OWNER
brm-interop-zs3-2-07_P0 (net/ixgbe0) 10.80.216.170	brm-interop-zs3-2-070
brm-interop-zs3-2-07_P1 (net/ixgbe1) 10.80.216.173	brm-interop-zs3-2-070
Untitled Interface (net/ixgbe2) 10.80.219.222	brm-interop-zs3-2-070
zfs/pool-1 3.20T	brm-interop-zs3-2-070

Active Resources

RESOURCE ▲	OWNER
zfs/pool-2	brm-interop-zs3-2-080

The following table describes characteristics of different cluster resource types.

Table 2-3 Cluster Resource Types

Resource	Icon	Omnipresent	Taken over on failure
SINGLETON	🔓	No	Yes
REPLICA	None	Yes	N/A
PRIVATE	🔒	No	No
SYMBIOTE	None	Same as parent type	Same as parent type

When a new resource is created, it is initially assigned to the controller on which it is being created. This ownership cannot be changed unless that controller is in the `AKCS_OWNER` state. Therefore, either create resources on the controller that should own them normally, or take over before changing resource ownership. In general, it is possible to destroy resources from either controller, although destroying storage pools that are exported is not possible. Best results will usually be obtained by destroying resources on the controller that currently controls them, regardless of which controller is the assigned owner.

Most configuration settings, including service properties, users, roles, identity mapping rules, SMB autohome rules, and iSCSI initiator definitions are replicated on both controllers automatically. It is never necessary to configure these settings on both controllers, regardless of the cluster state. If one appliance is down when the configuration change is made, the change will be replicated to the other appliance when that appliance rejoins the cluster on next boot, prior to providing any service. There are a small number of exceptions:

- Share and LUN definitions and options can be set only on the controller that has control of the underlying pool, regardless of the controller to which that pool is ordinarily assigned.

- The configuration of the Identity service (the appliance name and location) is not replicated.
- Names given to chassis are visible only on the controller on which they were assigned.
- Each network route is bound to a specific interface. If each controller is assigned an interface with an address in a particular subnet, and that subnet contains a router to which the appliances should direct traffic, a route must be created for each such interface, even if the same gateway address is used. This allows each route to become active individually as control of the underlying network resources shifts between the two controllers. For more information, see Clustering Considerations for Networking.
- SSH host keys are not replicated and are never shared. Therefore if no private administrative interface has been configured, you can expect key mismatches when attempting to log into the CLI using an address assigned to a node that has failed. The same limitations apply to the SSL certificates used to access the BUI.

The basic model is that common configuration is transparently replicated, and administrators will assign a collection of resources to each appliance controller. Those resource assignments form the binding of network addresses to storage resources that clients expect to see. Regardless of which appliance controls the collection of resources, clients are able to access the storage they require at the network locations they expect.

Related Topics

- Clustering Considerations for Storage
- Clustering Considerations for Networking

Cluster Takeover and Failback

Takeover enables service to continue or resume normally when a cluster controller fails or loses power.

Takeover of the cluster by one of the controllers is automatically attempted when the controller detects that its peer is absent (for example, shut down or rebooting). After takeover, the controller that performed the takeover owns all cluster resources and provides all services.

If both controllers fail or are powered off, then upon simultaneous startup, the appliance software performs an arbitration procedure to determine which controller will continue with takeover.

Takeover can also be performed manually, which can be useful for testing purposes.

The failback operation changes the cluster configuration from `OWNER-STRIPPED` (active-passive) to `CLUSTERED-CLUSTERED` (active-active). Failback never occurs automatically.

Failback usually is performed:

- When a controller is back online after a takeover.
- As the last step of configuring a cluster. See Upgrading a Standalone Appliance to a Clustered Configuration (BUI).

If `controller-b` in a cluster fails or loses power, then `controller-a` in that cluster takes over the resources that had been assigned to `controller-b`, and provides all cluster services. After `controller-b` is repaired and booted, an administrator performs the failback operation to return `controller-b` to production service.

When `controller-b` is repaired and booted, that controller:

- Rejoins the cluster, resynchronizing its view of all resources, their properties, and their ownership.
- Waits for an administrator to perform a failback operation.

While `controller-b` is waiting, `controller-a` continues to provide all services. `Controller-a` is in the Active (takeover completed) or `AKCS_OWNER` state, and `controller-b` is in the Ready (waiting for failback) or `AKCS_STRIPPED` state.

The failback operation returns `controller-b` to production service. Since the failure of `controller-b`, `controller-a` has been providing all services. The failback operation restores resources that were owned by `controller-b` prior to the failure back to `controller-b`. The failback operation exports from `controller-a` all resources that are assigned to `controller-b`, and `controller-b` imports these resources. After a successful failback, both `controller-a` and `controller-b` are in the Active or `CLUSTERED` state.

During failback, a pool that cannot be imported by `controller-b` because the pool is faulted will cause `controller-b` to reboot. The failback operation fails, and `controller-a` continues to provide all services.

When scheduling a failback operation, consider the following:

- Failback is disruptive to clients of the cluster.
- Delaying failback is equally or more disruptive if the single active controller fails before the failback is performed.

To minimize service downtime, data is not collected and statistics and datasets are not available during failback and takeover operations. Requests to suspend or resume statistics are delayed until failback and takeover operations have completed. Data collection automatically resumes after failback and takeover operations have completed.

Related Topics

- Clustered Controller States
- Estimating and Reducing Takeover Impact

Configuration Changes in a Clustered Environment

The majority of appliance configuration is represented as either service properties or share/LUN properties. While share and LUN properties are stored with the user data on the storage pool itself and thus are always accessible to the current owner of that storage resource, service configuration is stored within each controller. To ensure that both controllers provide coherent service, all service properties must be synchronized when a change occurs or a controller that was previously down rejoins with its peer. Since all services are represented by replica resources, this synchronization is performed automatically by the appliance software any time a property is changed on either controller.

It is therefore unnecessary and redundant for administrators to replicate configuration changes. Standard operating procedures should reflect this attribute and call for making changes to only one of the two controllers once initial cluster configuration has been completed. The process of initial cluster configuration will replicate all existing configuration onto the newly-configured peer.

The following are best practices for clustered configuration changes:

- Make all storage and network configuration changes on the controller that currently controls (or will control, if a new resource is being created) the underlying storage or network interface resources.

- Make all other changes on either controller, but not both. The controller that you specify as the *primary* controller should depend on which of the controllers is functioning and the number of storage pools that have been configured.

Oracle ZFS Storage Appliance has no mechanism for making independent changes to system configuration on each controller. This simplification alleviates the need for centralized configuration repositories. The controller that is currently operating is assumed to have the correct configuration, and its peer will be synchronized to it when booting. The peer will adopt a set of configuration parameters that are already in use by an existing production system and are therefore highly likely to be correct. Best practice is to ensure that a failed controller rejoins the cluster as soon as it is repaired.

Related Topics

- Clustering Considerations for Storage
- Clustering Considerations for Networking

Clustering Considerations for Storage

When sizing Oracle ZFS Storage Appliance for use in a cluster configuration, consider the following points:

- Whether all pools are owned by the same controller, or pools are split between the two controllers.
- Whether you want pools with no single point of failure (NSPF).

Assigning storage pool ownership - Perhaps the most important decision is whether all storage pools will be assigned ownership to the same controller, or split between the two controllers. There are several trade-offs to consider, as shown in Clustering Considerations for Storage.

Generally, pools should be configured on a single controller except when optimizing for throughput during nominal operation or when failed-over performance is not a consideration. The exact changes in performance characteristics when a controller is in the failed-over state will depend on the nature and size of the workloads. Generally, the closer a controller is to providing maximum performance on any particular axis, the greater the performance degradation along that axis when the workload is taken over by that controller's peer. Of course, in the multiple pool case, this degradation will apply to both workloads.

Read cache devices are located in the controller or disk shelf, depending on your configuration.

Read cache devices, located in a controller slot (internal L2ARC), do not follow data pools in takeover or failback situations. A read cache device is only active in a particular cluster node when the pool that is assigned to the read cache device is imported on the node where the device resides. Unless additional configuration steps are taken, read cache will not be available for a pool that has migrated due to a failover event. In order to enable a read cache device for a pool that is not owned by the cluster peer, take over the pool on the non-owning node, and then add storage and select the cache devices for configuration. Read cache devices in a cluster node should be configured as described in the Configuring Storage. Write-optimized log devices are located in the storage fabric and are always accessible to whichever controller has imported the pool.

If read cache devices are located in a disk shelf (external L2ARC), read cache is always available. During a failback or takeover operation, read cache remains sharable between controllers. In this case, read performance is sustained. For external read cache configuration details, see Disk Shelf Configurations in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

Configuring NSPF - A second important consideration for storage is the use of pool configurations with no single point of failure (NSPF). Since the use of clustering implies that the application places a high premium on availability, there is seldom a good reason to configure storage pools in a way that allows the failure of a single disk shelf to cause loss of availability. The downside to this approach is that NSPF configurations require a greater number of disk shelves than do configurations with a single point of failure. When the required capacity is very small, installation of enough disk shelves to provide for NSPF at the desired RAID level might not be economical.

The following table describes storage pool ownership for cluster configurations.

Table 2-4 Clustering Considerations for Storage Pools

Variable	Single Controller Pool Ownership	Multiple Pools Owned by Different Controllers
Total throughput (nominal operation)	Up to 50% of total CPU resources, 50% of DRAM, and 50% of total network connectivity can be used to provide service at any one time. This is straightforward: only a single controller is ever servicing client requests, so the other is idle.	All CPU and DRAM resources can be used to provide service at any one time. Up to 50% of all network connectivity can be used at any one time (dark network devices are required on each controller to support failover).
Total throughput (failed over)	No change in throughput relative to nominal operation.	100% of the surviving controller's resources will be used to provide service. Total throughput relative to nominal operation may range from approximately 40% to 100%, depending on utilization during nominal operation.
I/O latency	Internal read cache is not available during a failback or takeover operation, which can significantly increase latencies for read-heavy workloads that fit into available read cache. Latency of write operations is unaffected. With external read cache configurations (EL2ARC), read performance is unaffected. Read cache is shared between cluster peers during a failback or takeover operation, resulting in no read latency.	Internal read cache is not available during a failback or takeover operation, which can significantly increase latencies for read-heavy workloads that fit into available read cache. Latency of both read and write operations may be increased due to greater contention for controller resources. This is caused by running two workloads on the surviving controller instead of the usual one. When nominal workloads on each controller approach the controller's maximum capabilities, latencies in the failed-over state may be extremely high. With external read cache configurations (EL2ARC), read performance is unaffected. Read cache is shared between cluster peers during a failback or takeover operation, resulting in no read latency.
Storage flexibility	All available physical storage can be used by shares and LUNs.	Only the storage allocated to a particular pool can be used by that pool's shares and LUNs. Storage is not shared across pools, so if one pool fills up while the other has free space, some storage may be wasted.

Table 2-4 (Cont.) Clustering Considerations for Storage Pools

Variable	Single Controller Pool Ownership	Multiple Pools Owned by Different Controllers
Network connectivity	All network devices in each controller can be used while that controller is providing service.	Only half of all network devices in each controller can be used while that controller is providing service. Therefore each pool can be connected to only half as many physically disjoint networks.

Related Topics

Clustered Controller States

Clustering Considerations for Networking

Network device, datalink, and interface failures do not cause a clustered subsystem controller to fail. To protect against network failures inside or outside of the appliance, IPMP and/or LACP should be used. A comprehensive approach to availability requires the correct configuration of the network and a network-wide plan for redundancy.

Network interfaces can be configured as either singleton or private resources, provided they have a static IP configuration. Interfaces configured using DHCP must be private. Using DHCP in clusters is discouraged. When configured as a singleton resource, all datalinks and devices used to construct an interface can be active on only one controller at a time. Likewise, corresponding devices on each controller must be attached to the same networks in order for service to be provided in a failed-over state. See the example in the following figure.

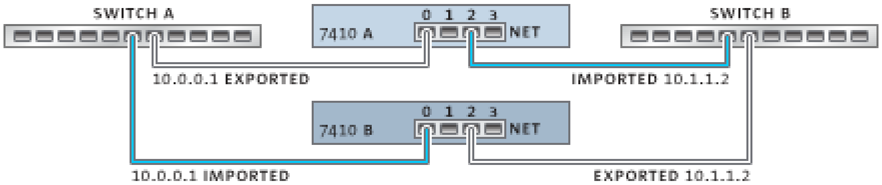

For a cluster to operate correctly when you construct network interfaces from devices and datalinks, it is essential that each singleton interface has a device using the same identifier and capabilities available on both controllers. Since device identifiers depend on the device type and the order in which they are first detected by the appliance, clustered controllers MUST have identical hardware installed. Each slot in both controllers must be populated with identical hardware and slots must be populated in the same order on both controllers. Your qualified Oracle reseller or service representative can assist in planning hardware upgrades that meet these requirements.

A route is always bound explicitly to a single network interface. Routes are represented within the resource manager as symbiotes. Routes can become active only when the interfaces to which they are bound are operational. Therefore, a route bound to an interface that is currently in standby mode (exported) has no effect until the interface is activated during the takeover process. This is important when two pools are configured and are made available to a common subnet. If a subnet is home to a route that is used by the appliances to reach one or more other networks, a separate route (for example, a second default route), must be configured and bound to each of the active and standby interfaces attached to that subnet.

Example:

- Interface e1000g3 is assigned to "alice" and e1000g4 is assigned to "bob".
- Each interface has an address in the 172.16.27.0/24 network and can be used to provide service to clients in the 172.16.64.0/22 network, reachable via 172.16.27.1.
- Two routes should be created to 172.16.64.0/22 via 172.16.27.1; one should be bound to e1000g3 and the other to e1000g4.

It is a good idea to assign each clustered controller an IP address used only for administration (most likely on a dedicated management network) and to designate the interface as a private resource. This ensures that it is possible to reach a functioning controller from the management network even if it is in an `AKCS_STRIPPED` state and awaiting failback. This is important if services such as LDAP and Active Directory are in use and require access to other network resources when the controller is not providing service. If this is not practical, the service processor should be attached to a reliable network and/or serial terminal concentrator so that the controller can be managed using the system console.

If neither of these actions is taken, it is impossible to manage or monitor a newly-booted controller until failback is completed. You might want to monitor or manage the controller that is providing service for a particular storage pool. This is likely to be useful when you want to modify some aspect of the storage itself such as a share property, or create a new LUN. This can be done by using one of the service interfaces to perform administrative tasks or by allocating a separate singleton interface to be used only for managing the pool to which it is matched. In either case, the interface should be assigned to the same controller as the pool it is used to manage.

Impact to NFSv4.1 clients - Certain networking changes in a cluster configuration can adversely affect the servicing of requests for NFSv4.1 clients. If the relationship between an IP address and its owner changes, then the best practice is to remount the filesystems from the client. Unlike NFSv4.0, NFSv4.1 protocol enables client connections over multiple IP addresses to be associated with the same NFSv4.1 protocol lease. When the relationship between an IP address and its owner changes, the group of IP addresses that failover together is no longer the same, forcing the client to re-establish the lease relationships by remounting the filesystems.

Related Topics

- Private Local IP Interfaces
- Clustering Considerations for InfiniBand
- Clustered Controller States

Private Local IP Interfaces

Use the following guidelines when creating private local IP interfaces:

- Creating an IP interface with the same name as a private IP interface on a cluster peer results in the local creation of a private IP interface.
- Datalinks in use by the peer's private interfaces cannot be deleted, and the delete button is grayed out.
- IP interfaces that belong to an IPMP group must all be of the same type and belong to the same controller. To create an IPMP group you must use either all singleton or all private IP interfaces and your cluster node must be the owner of these interfaces.
- The IPMP group type is set only at creation, and is determined by the type of underlying links.

- IP interfaces that belong to IPMP groups do not appear on the **Cluster: Resources** page because IP interface ownership cannot be modified independently of the IPMP group ownership.
- Private IPMP groups do not appear in the **Cluster: Resources** page because this type or ownership cannot be modified.

Related Topics

Clustering Considerations for Networking

Clustering Considerations for InfiniBand

> **Note:**
>
> Oracle ZFS Storage ZS11-2 and ZS9-2 controllers do not support InfiniBand.

Like a network built on top of Ethernet devices, an InfiniBand network needs to be part of a redundant fabric topology in order to guard against network failures inside and outside of the appliance. The network topology should include IPMP to protect against network failures at the link level with a broader plan for redundancy for HCAs, switches, and subnet managers.

To ensure proper cluster configuration, each controller must be populated with identical HCAs in identical slots. Furthermore, each corresponding HCA port must be configured into the same partition (`pkey`) on the subnet manager with identical membership privileges and attached to the same network. To reduce complexity and ensure proper redundancy, it is recommended that each port belong to only one partition in the InfiniBand sub-network. Network interfaces may be configured as either singleton or private resources, provided they have static IP configuration. When configured as a singleton resource, all of the InfiniBand partition datalinks and devices used to construct an interface may be active on only one controller at any given time. An example of this is shown in the following figure.

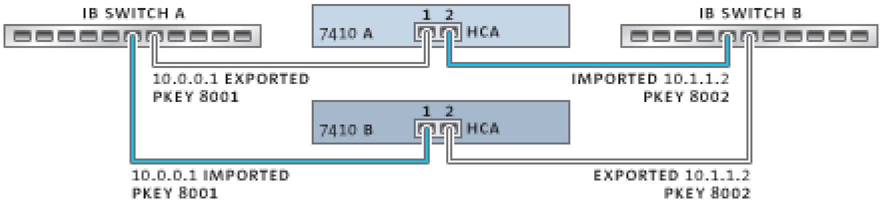

Changes to partition membership for corresponding ports must happen at the same time. Your qualified Oracle reseller or service representative can assist in planning hardware upgrades that will meet these requirements.

The following figure shows cluster configuration for subnet manager redundancy. Greater redundancy is achieved by connecting two dual-port HCAs to a redundant pair of server switches.

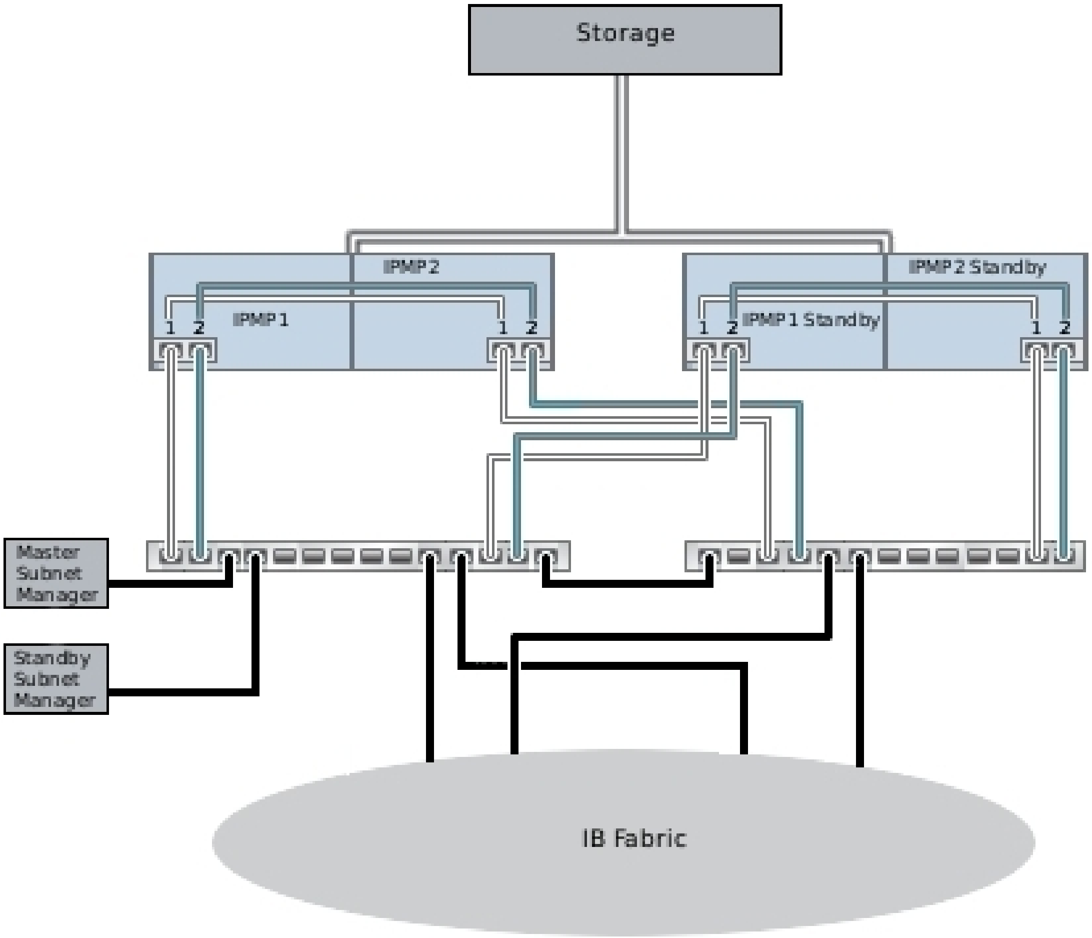

Related Topics

Clustering Considerations for Networking

Preventing Split-Brain Conditions

A common failure mode in clustered systems is known as *split-brain*. In this condition, each of the clustered controllers behaves as if its peer has failed and attempts takeover. The most common cause of this condition is failure of the communication medium shared by the controllers. In Oracle ZFS Storage Appliance, the shared communication medium is the cluster I/O links. However, Oracle ZFS Storage Appliance cluster I/O links have built-in link redundancy: For Oracle ZFS Storage ZS11-2 and Oracle ZFS Storage ZS9-2 controllers, only a single cluster I/O Ethernet link is required to avoid triggering takeover. For all other controllers, only a single cluster I/O serial link is required to avoid triggering takeover.

The appliance software performs an arbitration procedure to determine which controller should continue with takeover.

The Oracle ZFS Storage Appliance clustering solution is designed to ensure that there is no single point of failure, and to protect both data and availability against failure. Most failures can be prevented by installing the hardware properly and employing cluster setup and management best practices. Ensure the following:

- ll cluster I/O links (two for an Oracle ZFS Storage ZS11-2 or an Oracle ZFS Storage ZS9-2 controller, three for all other controllers) are connected and functional as shown in Cluster Configuration BUI View and Checking Cluster Link Status (CLI).
- All storage cabling is connected as shown in the setup documentation that was delivered with your appliances.

  It is particularly important that two paths are detected to each disk shelf as shown in the following figure before placing the cluster into production and at all times afterward, with the exception of temporary cabling changes to support capacity increases or replacement of faulty components. Use alerts to monitor the state of cluster interconnect links and disk shelf paths and correct any failures promptly. Ensuring that proper connectivity is maintained will protect both availability and data integrity if a hardware or software component fails.

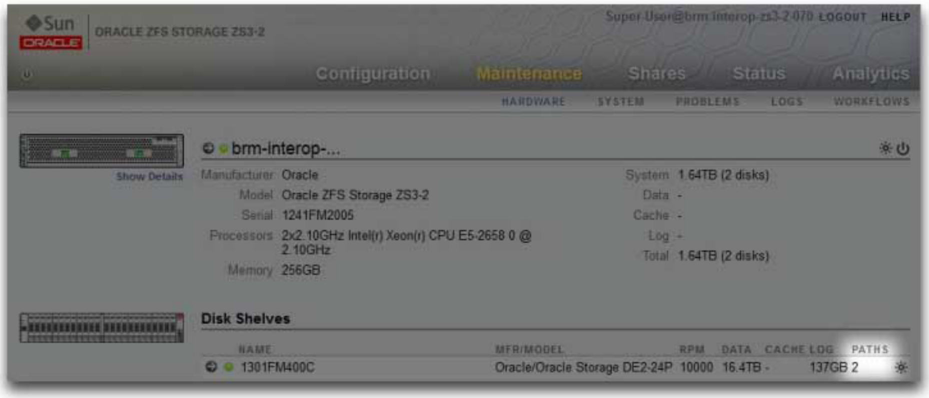

brm-interop-...

Manufacturer Oracle
Model Oracle ZFS Storage ZS3-2
Serial 1241FM2005
Processors 2x2.10GHz Intel(r) Xeon(r) CPU E5-2658 0 @ 2.10GHz
Memory 256GB

System 1.64TB (2 disks)
Data -
Cache -
Log -
Total 1.64TB (2 disks)

Disk Shelves

NAME	MFR/MODEL	RPM	DATA	CACHE	LOG	PATHS
1301FM400C	Oracle/Oracle Storage DE2-24P	10000	16.4TB	-	137GB	2

Related Topics

Clustered Controller States

Estimating and Reducing Takeover Impact

During takeover and failback, there is an interval during which access to storage cannot be provided to clients. The length of this interval varies by configuration, and the exact effects on clients depends on the protocols they are using to access data. Understanding and mitigating these effects can make the difference between a successful cluster deployment and a costly failure at the worst possible time.

NFS (all versions) clients typically hide outages from application software, causing I/O operations to be delayed while a server is unavailable. NFSv2 and NFSv3 are stateless protocols that recover almost immediately upon service restoration. NFSv4.0 and NFSv4.1 incorporate a client grace period at startup, during which I/O typically cannot be performed. The duration of this grace period can be tuned in Oracle ZFS Storage Appliance. Reducing this grace period will reduce the apparent impact of takeover and/or failback. For planned outages, the appliance provides grace-less recovery for NFSv4.0 and NFSv4.1 clients, which avoids the grace period delay. For more information about grace-less recovery, see the description of the grace period property in NFS Service Properties. The following figure shows the NFS grace period property.

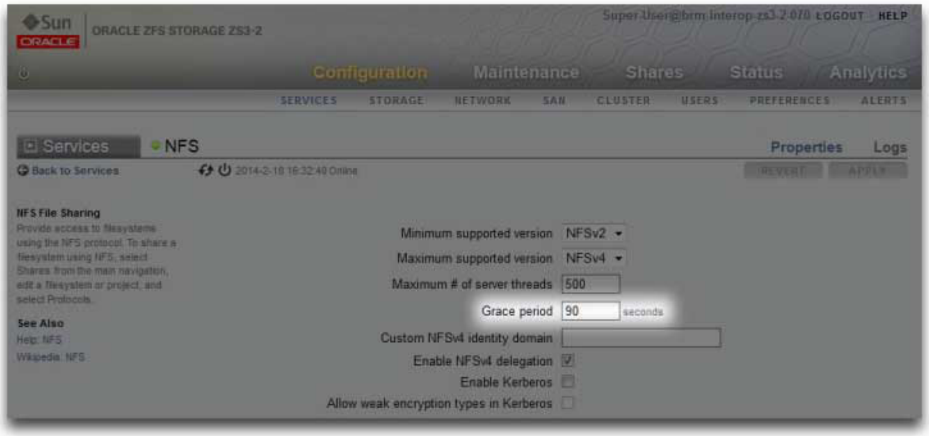

2014-2-18 16:32:49 Online

NFS File Sharing

Provide access to filesystems using the NFS protocol. To share a filesystem using NFS, select Shares from the main navigation, edit a filesystem or project, and select Protocols.

See Also

Help: NFS

Wikipedia: NFS

Minimum supported version: NFSv2

Maximum supported version: NFSv4

Maximum # of server threads: 500

Grace period: 90 seconds

Custom NFSv4 identity domain:

Enable NFSv4 delegation ☑

Enable Kerberos □

Allow weak encryption types in Kerberos □

iSCSI behavior during service interruptions is initiator-dependent, but initiators will typically recover if service is restored within a client-specific timeout period. Check your initiator's documentation for additional details. The iSCSI target will typically be able to provide service as soon as takeover is complete, with no additional delays.

SMB, FTP, and HTTP/WebDAV are connection-oriented protocols. Because the session states associated with these services cannot be transferred along with the underlying storage and network connectivity, all clients that use one of these protocols will be disconnected during a takeover or failback, and must reconnect after the operation completes.

While several factors affect takeover time and failback time, in most configurations these times will be dominated by the time required to import the diskset resources. Typical import times for each diskset range from 15 to 20 seconds, linear in the number of disksets. A diskset consists of one half of one disk shelf, provided the disk bays in that half-disk shelf have been populated and allocated to a storage pool. Unallocated disks and empty disk bays have no effect on takeover time. The time taken to import diskset resources is not affected by any parameters that can be tuned or altered by administrators. Therefore, do one of the following if you are planning a clustered deployment:

- Limit installed storage so that clients can tolerate the related takeover times
- Adjust client-side timeout values above the maximum expected takeover time

Note that while diskset import usually comprises the bulk of takeover time, it is not the only factor. During the pool import process, any intent log records must be replayed, and each share and LUN must be shared via the appropriate services. The amount of time required to perform these activities for a single share or LUN is very small - on the order of tens of milliseconds - but with very large share counts this can contribute significantly to takeover times. Keeping the number of shares relatively small - a few thousand or fewer - can reduce these times considerably.

Failback time is normally greater than takeover time for any given configuration. This is because failback is a two-step operation: First, the source appliance exports all resources of which it is not the assigned owner, then the target appliance performs the standard takeover procedure on its own assigned resources only. Therefore it will always take longer to failback from controller A to controller B than it will take for controller A to take over from controller B in case of failure. This additional failback time is much less dependent upon the number of disksets being exported than is the takeover time, so keeping the number of shares and LUNs small can have a greater impact on failback than on takeover. It is also important to keep in mind that failback is always initiated by an administrator, so the longer service interruption it causes can be scheduled for a time when it will cause the lowest level of business disruption.

> **Note:**
>
> Estimated durations cited in this section can vary per software/firmware version. It is important to test the takeover operation and its exact impact on client applications prior to deploying a clustered configuration in a production environment.

Related Topics

- Cluster Takeover and Failback
- Clustered Controller States

Network Configuration

The network configuration features enable you to create a variety of advanced network configurations using your physical network ports (including link aggregations), virtual NICs (VNICs), virtual LANs (VLANs), and multipathing groups. You can define any number of IPv4 and IPv6 addresses for these objects to connect to the various data services on the system.

The network configuration of an appliance has the following components. To configure networking, build datalinks on devices and build interfaces on datalinks. Select an object to view its relationship to other objects.

- **Devices** - Network devices represent the physical network connections or IP over InfiniBand (IPoIB) partitions that correspond to your physical network. Network devices are created by the system and have no configurable settings.

  Devices that represent physical network interface card (NIC) ports on the controller are typically labeled igb0, igb1, igb2, and igb3. For example, device igb0 represents port NET-0 and device igb1 represents port NET-1. One NIC port per controller should be configured as a *management interface*.

  Devices that represent physical InfiniBand ports on the controller are typically labeled ibp0, ibp1, ibp2, and ibp3.

- **Datalinks** - Datalinks are Layer 2 objects for sending and receiving packets for specific network devices. Use datalinks to apply settings such as LACP to network devices. Datalinks can correspond 1:1 with a device or IB partition, or you can define aggregation, VLAN, and VNIC datalinks composed of other devices and datalinks. Datalinks are required to complete network configuration, even if the datalinks do not apply specific settings to network devices. See Configuring Network Datalinks.

- **Interfaces** - Interfaces are Layer 3 objects for IP, configuring IP addresses and other properties for datalinks. Each IP interface is associated with a single datalink, or is defined to be an IP network multipathing (IPMP) group comprising a pool of datalinks, which allows automatic migration of IP addresses from failed to working datalinks. See Configuring Network Interfaces.

  The following example shows configuration for a single IP address on a single port:

  - Device - igb0
  - Datalink - datalink1
  - Interface - *hostname* (*IP-address*/*mask*)

  The following example shows configuration for a 3-way link aggregation:

  - Devices - igb1, igb2, igb3
  - Datalink - aggr1 (LACP aggregation)

  - Interface - *hostname* (*IP-address*/*mask*)
- **Routing** - IP routing configuration controls how the system will direct IP packets. See Configuring Network Routing.

To configure the network for the appliance, use the following sections:

- Configuring Network Datalinks
- Configuring Network Interfaces
- Configuring Network IP Multipathing (IPMP)
- Configuring Network Performance and Availability
- Configuring Network Routing

Using the BUI to Configure Networking

To ensure that you maintain your browser connection while you make network configuration changes, assign a particular IP address and network device for administrative use as described in Configuring Management Interfaces.

The **Configuration: Network: Configuration** page lists devices, datalinks, and interfaces. The **Devices** list shows link status on the right, and shows an icon next to the device name that indicates the state of the network port. If one of these icons shows that the device is down (see the following table), check that the device is plugged into the network properly.

Some configuration can be performed by dragging and dropping a device entry onto the datalinks column or a datalink entry onto the interfaces column. This can be helpful for complex configurations. Valid moves are highlighted.

The **Configuration: Network: Addresses** page lists the interface name, IP address, and hostname for each datalink.

The **Configuration: Network: Routing** page provides the following:

- Shows the multihoming model: Loose, Adaptive, or Strict
- Shows the destination, gateway, family, type, status, and interface for each route
- Provides the ability to add a static route

The following table shows icons that are used in network configuration.

Table 2-5 Network Configuration Icons

Icon	Description
	Connected network port
	Connected network port with I/O activity
	Disconnected network port (link down, cable problem)
	Active InfiniBand port
	Active InfiniBand port with I/O activity
	Inactive InfiniBand port (down, init, or arm state)

Table 2-5 (Cont.) Network Configuration Icons

Icon	Description
▟	InfiniBand partition device is up
●	InfiniBand partition device is down (subnet manager problem)
<••>	Network datalink
<∘∘>	Network datalink VLAN or VNIC
{:::}	Network datalink aggregation
{∘∘∘}	Network datalink aggregation VLAN or VNIC
🔗	Network datalink IB partition
▟	Interface is being used to send and receive packets (either up or degraded)
●	Interface has been disabled by the user
●	Interface is offline (owned by the cluster peer)
●	Interface has failed or has been configured with a duplicate IP address

Using the CLI to Configure Networking

Network configuration is under the `configuration net` path, which has subcommands for `devices`, `datalinks`, `interfaces`, and `routing`.

Use the `available` command to see what values can be assigned to the `links` property when creating a network component, as shown in the following example:

```
hostname:configuration net interfaces> ip
hostname:configuration net interfaces ip (uncommitted)> available
igb0,vnic1
```

Configuring Network Datalinks

Network datalinks manage devices, and are used by interfaces. A datalink can be a VLAN, VNIC, IB Partition, or an LACP aggregation of devices.

- **Virtual LAN (VLAN)** - VLANs are used to improve local network security and isolation. VLANs are recommended for administering the appliance; otherwise, use VNICs.
- **Virtual Network Interface Card (VNIC)** - VNICs allow single or aggregated Ethernet datalinks to be split into multiple virtual (Ethernet) datalinks. VNICs can be optionally tagged with VLAN IDs, and can allow physical network port sharing in a cluster as described in Clustering Considerations for Networking.
- **InfiniBand (IB) partitions** - InfiniBand partitions connect to logically isolated IB fabric domains.
- **Link Aggregation Control Protocol (LACP)** - LACP is used to bundle multiple network devices such that they behave as one. This improves performance by increasing bandwidth and improves reliability by protecting from network port failure. The appliance must be connected to a switch that supports LACP and has enabled LACP for those ports.

> **Note:**
>
> VNIC-based and VLAN-based datalinks cannot share the same VLAN ID. The IEEE 802.3ad Link Aggregation standard does not explicitly support aggregations across multiple switches, but some vendors provide multi-switch support via proprietary extensions. If a switch configured with those extensions conforms to the IEEE standard and the extensions are transparent to the end-nodes, its use is supported with the appliance. If an issue is encountered, Oracle Support might require the issue to be reproduced on a single-switch configuration.

The following table describes settings that are available to configure datalinks.

Table 2-6 Datalink Properties

BUI Property	CLI Property	Description
Name	`label`	Use the defined custom name. For example: "`internal`", "`external`", "`adminnet`", and so on. The default value is "`Untitled Datalink`".
Speed	`speed`	Use the defined speed. Valid values are `auto, 10, 100, 1000` and `10000`, representing autonegotiation, forced 10Mbit/sec, forced 100Mbit/sec, forced 1Gbit/sec and forced 10Gbit/sec. Speed and duplex must be either both forced to specific values or both set to autonegotiate. Not all networking devices support forcing to all possible speed/duplex combinations. Disabling autonegotiation is strongly discouraged. However, if the switch has autonegotiation disabled, it might be necessary to force speed and duplex to ensure the datalink runs at the expected speed and duplex.
Duplex	`duplex`	Use the defined transmission direction. Valid CLI values are `auto, half`, and `full`, representing autonegotiation, half-duplex and full-duplex respectively. Speed and duplex must be either both forced to specific values or both set to autonegotiate. See more information in the description of the speed property earlier.
VLAN	`class=vlan`	Use VLAN headers.
VLAN ID	`id`	Use the defined VLAN identifier; optional for VNICs.
VNIC	`class=vnic`	Use a VNIC.
MTU	`mtu`	Use the defined maximum transmission unit (MTU) size. The default MTU is `1500` bytes. Specify a lower MTU (minimum 1280) to leave packet headroom (for example, for tunneling protocols). Specify a larger MTU (maximum 9000) to improve network performance. All systems and switches on the same LAN must be configured with the chosen MTU. After the MTU value is set and the new network configuration is committed to the system, you can return to the network screen and view the datalink status to see the exact MTU value in bytes that was selected. Note that a VLAN or VNIC cannot be configured with an MTU value larger than that of the underlying datalink.
LACP Aggregation	`aggregation`	Use multiple network device LACP aggregation.

Table 2-6 (Cont.) Datalink Properties

BUI Property	CLI Property	Description
LACP Policy	`policy`	Use the defined LACP policy for selecting an outbound port. `L2` hashes the source and destination MAC address; `L3` uses the source and destination IP address; `L4` uses the source and destination transport level port.
LACP Mode	`mode`	Use the defined LACP communication mode. `Active` mode will send and receive LACP messages to negotiate connections and monitor the link status. `Passive` mode will listen for LACP messages only. `Off` mode will use the aggregated link but not detect link failure or switch configuration changes. Some network switch configurations, including Cisco EtherChannel, do not use the LACP protocol: The LACP mode should be set to `off` when using non-LACP aggregation in your network.
LACP Timer	`timer`	Use the defined interval between LACP messages for Active mode.
IB Partition	`partition`	Use IB Partitions.
Partition Key	`pkey`	Use the partition (fabric domain) in which the underlying port device is a member. The partition key is found on and configured by the subnet manager. The partition key can be defined before configuring the subnet manager but the datalink will remain "down" until the subnet partition has been configured with the port GUID as a member. Keep partition membership for HCA ports consistent with IPMP Configuration and Appliance Cluster Configuration rules on the subnet manager.
IB Link Mode	`linkmode`	Use the defined IB link mode. IPoIB provides two link modes: `Connected` (the default) and `Unreliable Datagram`. `Connected` mode is recommended over `Unreliable Datagram`. Use `Unreliable Datagram` only when technically required. `Connected` mode provides higher throughput than `Unreliable Datagram` mode. `Connected` mode uses an MTU of 65520. `Unreliable Datagram` mode uses an MTU of 2044. `Connected` mode uses IB queue pairs and dedicates a local queue pair to communicate with a dedicated remote queue pair. `Unreliable Datagram` allows a local queue pair to communicate with multiple other queue pairs on any host, and messages are communicated unacknowledged at the IB layer.
Large Receive Offload (LRO)	`lro`	Use large receive offload (LRO) (`on`). Default is `off`. LRO merges successive incoming packets into a single packet before the packets are delivered to the IP layer. Enabling this feature might provide performance benefits when using smaller MTU sizes. LRO does not apply to IPv6 interfaces that are built on top of datalinks with LRO enabled.

Configuring a Network Datalink (BUI)

Use this procedure to configure datalinks, which can then be used to configure interfaces.

1. From the **Configuration** menu, select **Network**, then **Configuration**.

2. Click the add icon ⊕ located next to **Datalinks**.
3. Optional: In the **Network Datalink** dialog box, type a **Name** for the datalink.

   By default the name of the datalink is set to `Untitled Datalink`.
4. Optional: Set a **Custom MTU**.

   For the `Max Transmission Unit (MTU)` property, specify the following:

   a. Click the **Custom** button.

   b. Type `9000` in the text field.
5. Optional: Set the **Link Speed, Link Duplex**, and **Flow Control**.

   These settings are available if you did not select either the **VLAN** or **VNIC** check box at the top of the dialog box.
6. Optional: Set a **VLAN ID**.

   This field is available if you selected the **VLAN** or **VNIC** check box at the top of the dialog box.
7. Choose a device from the **Devices** list.
8. Click **APPLY**.

   The datalink appears in the **Datalinks** list.

Next Steps

- To modify a datalink, double-click the datalink or hover over the datalink and click the edit icon ✎ , make the changes, and click the **APPLY** button.
- To delete a datalink, hover over the datalink, click the delete icon 🗑 , and confirm that you want to delete the datalink.

Related Topics

Configuring a Management Interface (BUI)

Configuring a Network Datalink (CLI)

Use this procedure to configure datalinks, which can then be used to configure interfaces.

1. Go to `configuration net datalinks` and enter the `device` command.

   ```
   hostname:configuration net datalinks> device
   ```
2. Optional: Set the name of the datalink.

   By default the `label` is set to `Untitled Datalink`.

   ```
   hostname:configuration net datalinks device (uncommitted)> set label=datalink2
                             label = datalink2 (uncommitted)
   ```
3. Set a device for this datalink.

   Set the value of the `links` property to the name of an existing device. Use the `available` command to see a list of existing devices.

   ```
   hostname:configuration net datalinks device (uncommitted)> available
   igb0,igb1
   hostname:configuration net datalinks device (uncommitted)> set links=igb1
                             links = igb1 (uncommitted)
   ```
4. Optional: Set the Max Transmission Unit (MTU).

By default, mtu is set to 1500.

5. Commit the changes.

```
hostname:configuration net datalinks device (uncommitted)> commit
hostname:configuration net datalinks> show
Datalinks:

DATALINK       CLASS        LINKS        STATE   ID      LABEL
igb0           device       igb0         up      -       datalink1
igb1           device       igb1         up      -       datalink2
```

Related Topics

Configuring a Management Interface (CLI)

Configuring Network Interfaces

Interfaces configure IP addresses via datalinks. Interfaces support the following features:

- IPv4 and IPv6 protocols.
- IPMP - IP network multipathing, to improve network reliability by allowing IP addresses to automatically migrate from failed to working datalinks.

The following settings are available for interfaces.

Table 2-7 Interface Properties

BUI Property	CLI Property	Description
Name	label	Custom name for the interface.
Enable Interface	enable	Enable this interface to be used for IP traffic. If an interface is disabled, the appliance will no longer send or receive IP traffic over it, or make use of any IP addresses configured on it. At present, disabling an active IP interface in an IPMP group will not trigger activation of a standby interface.
Allow Administration	allow	Allow connections to the appliance administration BUI or CLI over this interface. If your network environment included a separate administration network, this could be enabled for the administration network only to improve security. See Configuring a Management Interface - BUI, CLI.
IPv4 Configure with	v4addrs or v4dhcp	Select Static Address List (or enter values for v4addrs) or select DHCP (or set v4dhcp to true) for dynamically requested.
IPv4 Address/Mask	v4addrs	One or more IPv4 addresses in CIDR notation (192.168.1.1/24).
IPv6 Configure with	v6addrs or v6dhcp	Select Static Address List (or enter values for v6addrs) or select IPv6 AutoConfiguration (or set v6dhcp to true) to use automatically generated link-local address (and site-local if an IPv6 router responds).
IPv6 Address/Mask	v6addrs	One or more IPv6 addresses in CIDR notation (1080::8:800:200C:417A/32).

Table 2-7 (Cont.) Interface Properties

BUI Property	CLI Property	Description
Directly Reachable Network(s)	`v4directnets` or `v6directnets`	Directly reachable subnet(s), expressed as an IP address and mask in CIDR notation, that the local IP address is not a member of, but to which the datalink of its interface is physically connected. This improves scalability by conserving IP addresses, and could ease traffic congestion through core switches and routers.
IP MultiPathing Group	`links`	Configure IP network multipathing, where a pool of datalinks can be used for redundancy.

Creating a Single Port Interface (BUI)

Use this procedure to configure a single port network interface.

Before You Begin

At least one datalink must already exist. To create a datalink, see Configuring a Network Datalink (BUI).

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Drag a datalink from the **Datalinks** list to the **Interfaces** list, or click the **Interface** add item icon ⊕ .
3. In the **Network Interface** dialog box, set the desired properties.
   - **Name -** Type a name for the interface. You might want the interface name to indicate that this is a management interface.
   - **Enable Interface -** Select this check box to enable the interface.
   - **Use IPv4 Protocol** or **Use IPv6 Protocol -** Select a protocol, its type of address, and enter one or more IP addresses in CIDR notation.

   For the **Allow Administration** property, see Configuring a Management Interface (BUI).
4. Choose a datalink from the **Datalinks** list.
5. Click **APPLY** in the **Network Interface** dialog box.

   The interface appears in the **Interfaces** list.
6. Click **APPLY** in the upper-right corner of the **Configuration: Network: Configuration** page.

   The running appliance network configuration does not change until you perform this step.

Next Steps

- Repeat this procedure for as many IPs as needed. Ownership of an interface is defined by the node that the interface was created on. If this is a clustered configuration, create interfaces on the second controller.
- To modify this interface, double-click the interface or hover over the interface and click the edit icon ✎ , make the changes, and click the **APPLY** button.
- To delete this interface, hover over the interface, click the delete icon 🗑 , and confirm that you want to delete this interface.

> **Note:**
>
> When an interface is deleted, all routes associated with the interface are also removed.

Related Topics

- Configuring a Network Datalink (BUI)
- Configuring a Management Interface (BUI)

Unlocking Data Interfaces (BUI)

Use this procedure to unlock data interfaces.

1. From the **Configuration** menu, select **Cluster**.
2. Choose a data interface from the **Resource** list.
3. If the color of the lock icon is black 🔒, click the icon to unlock the interface.

   When the interface is unlocked, the lock icon is gray 🔒.
4. Click **APPLY** in the upper-right corner of the page.

Unlocking Data Interfaces (CLI)

Use this procedure to unlock data interfaces.

1. Go to `configuration cluster resources` and enter the `show` command.
2. Ensure that all data interfaces have the type property set to `singleton`.
3. If a data interface has type set to `private`, select the data interface and set the type to `singleton`.

   ```
   hostname:configuration cluster resources zfs/data1> set type=singleton
                               type = singleton (uncommitted)
   hostname:configuration cluster resources zfs/data1> commit
   ```

Creating an InfiniBand Partition Datalink and Interface (BUI)

Use this procedure to create an InfiniBand partition datalink and interface.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Click the **Datalinks** add item icon ⊕.
3. Optionally, set a name.
4. Click the **IB Partition** check box.
5. Choose a device from the **Partition Devices** list.
6. Enter a four-digit hexadecimal number for the partition key, which must match what was configured on the InfiniBand subnet manager.
7. Choose link mode from the drop-down menu.
8. Click **APPLY**.

   The new partition datalink will appear in the **Datalinks** list.
9. Click the Interface add item icon ⊕.

10. Set desired properties, and choose the datalink previously created.
11. Click **APPLY**.

    The interface will appear in the **Interfaces** list.
12. The running appliance network configuration has not yet changed. When you are finished configuring interfaces, click **APPLY** at the top to commit the configuration.

Creating a VNIC Without a VLAN ID for Clustered Controllers (BUI)

This example is for an active-active configuration with half of the network ports on standby. This task creates an IP interface over a device datalink and assigns it to a head. A VNIC is built on top of the same datalink, and an IP interface is configured on top of the VNIC and assigned to the other head. Configuring one instead of multiple VNICs over a given datalink ensures peak performance. Traffic flows over the cable associated with the underlying active port on one head, as well as the underlying standby port on the other head. Thus, the otherwise idle standby port can be used with VNICs.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. When the cluster is in state `AKCS_CLUSTERED`, click the **Datalinks** add item icon ⊕ .
3. Optional: Set a name and MTU value.
4. Choose a device from the **Devices** list, and click **APPLY**.

   The datalink appears in the **Datalinks** list.
5. Click the **Interface** add item icon ⊕ .
6. Set desired properties, choose the datalink previously created, and click **APPLY**.

   The interface appears in the **Interfaces** list.
7. Click the **Datalinks** add item icon ⊕ .
8. Select the **VNIC** check box, optionally set the name and MTU (equal to or less than the value in step 2), and click **APPLY**.

   The new VNIC datalink appears in the **Datalinks** list.
9. Click the **Interface** add item icon ⊕ .
10. Set desired properties, choose the VNIC datalink previously created, and click **APPLY**.

    The interface appears in the **Interfaces** list.
11. The running appliance network configuration has not yet changed. When you are finished configuring interfaces, click **APPLY** at the top to commit the configuration.
12. Click the **Cluster** tab.

    The two newly created interfaces appear in the **Resource** section with default owners.
13. Use the **Owner** pull-down list to assign one of the two interfaces to the other controller, and click **APPLY**.

Creating VNICs with the Same VLAN ID for Clustered Controllers (BUI)

This example is for an active-active configuration with half of the network ports on standby. This task creates two VNICs with identical VLAN IDs on top of the same device datalink. Each VNIC is configured with an interface, and each interface is assigned to a different head. Traffic flows over the cable associated with the underlying active port on one head, as well as the underlying standby port on the other head. Thus, the otherwise idle standby port can be used with VNICs.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. When the cluster is in state `AKCS_CLUSTERED`, click the **Datalinks** add item icon ⊕ .
3. Select the **VNIC** check box, optionally set the name and MTU, set the VLAN ID, choose a device from the **Devices** list, and click **APPLY**.

   The new VNIC datalink appears in the **Datalinks** list.
4. Click the **Interface** add item icon ⊕ .
5. Set desired properties, choose the **VNIC** datalink previously created, and click **APPLY**.

   The interface appears in the **Interfaces** list.
6. Create another VNIC as described in steps 2 and 3 with the same `Device` and `VLAN ID`, and create an interface for it as described in steps 4 and 5.
7. The running appliance network configuration has not yet changed. When you are finished configuring interfaces, click **APPLY** at the top to commit the configuration.
8. Click the **Cluster** tab.

   The two newly created interfaces appear in the **Resource** section with default owners.
9. Use the **Owner** pull-down list to assign one of the two interfaces to the other controller, and click **APPLY**.

Configuring Management Interfaces

If you did not set a management interface during initial configuration, configure a network interface card (NIC) port as a management interface. A management interface is a network interface with administrative access.

All standalone controllers should have at least one NIC port configured as a management interface. All cluster installations should have at least one NIC port on each controller configured as a management interface.

Set the following configuration for clustered controllers:

- The NIC instance number must be unique on each controller.
- The management interfaces should be locked.

A management interface enables BUI connections on port 215 and CLI connections on `ssh` port 22.

To configure management interfaces and lock cluster management interfaces, use the following procedures:

- Configuring a Management Interface - BUI, CLI
- Locking a Cluster Management Interface - BUI, CLI

Configuring a Management Interface (BUI)

Use this procedure to configure a management, or administrative, interface.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Select a datalink.

   You can modify an existing datalink that has the VNIC property set, or you can create a new datalink for this interface.

To create a new datalink, perform the following steps. See also Configuring a Network Datalink (BUI).

a. Click the add icon ⊕ next to **Datalinks**.

b. Select the check box for **VNIC**.

c. In the **Name** field, type a name for the datalink.

   You might want the datalink name to indicate that this is for a management interface.

d. In the **Network Datalink** dialog box, click **APPLY**.

   The datalink appears in the **Datalinks** list.

3. Drag this datalink from the **Datalinks** list to the **Interfaces** list, or click the **Interface** add item icon ⊕ .
4. In the **Network Interface** dialog box, set the following properties.
   - **Name** - Type a name for the interface.
   - **Enable Interface** - Select this check box to enable the interface.
   - **Use IPv4 Protocol** or **Use IPv6 Protocol** - Select a protocol, its type of address, and enter one or more IP addresses in CIDR notation.
5. Select the check box for **Allow Administration**.

   Select this check box to set this interface as a management interface. A management interface enables BUI connections on port 215 and CLI connections on `ssh` port 22.
6. From the **Datalinks** list, choose a datalink.
7. In the **Network Interface** dialog box, click **APPLY**.

   The interface appears in the **Interfaces** list.
8. In the upper right corner of the **Configuration: Network: Configuration** page, click **APPLY**.

   The running appliance network configuration does not change until you perform this step.
9. In the **Update Default Route** dialog box, set the following properties.
   - **Default Gateway** - This is the default router IP address.
   - **Interface** - Select the datalink that you assigned to the first management interface.

> **Note:**
>
> This step is strongly recommended because setting a route enables communication with the appliance via the BUI and CLI. Without a route, the only means of communication with the appliance is through an Oracle ILOM connection to the SP.

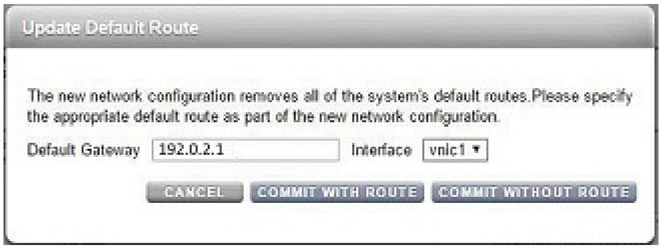

Update Default Route

The new network configuration removes all of the system's default routes.Please specify the appropriate default route as part of the new network configuration.

Default Gateway 192.0.2.1 Interface vnic1 ▾

10. Click **COMMIT WITH ROUTE**.
11. If this is a clustered configuration, lock the interface.

    See Locking a Cluster Management Interface (BUI). This step is optional but strongly recommended.

Next Steps

- If this is a clustered configuration, repeat this procedure on the second controller. Ownership of an interface is defined by the node that the interface was created on. If an interface is needed for the second controller in a cluster, create the interface on that second controller.
- To modify this interface, double-click the interface, or hover over the interface and click the edit icon ✎ and make the changes, then click the **APPLY** button.
- To delete this interface, hover over the interface, click the delete icon 🗑 and confirm that you want to delete this interface.

> **Note:**
>
> When an interface is deleted, all routes associated with the interface are also removed.

Related Topics

- Configuring a Network Datalink (BUI)
- Locking a Cluster Management Interface (BUI)

Configuring a Management Interface (CLI)

Use this procedure to configure a management, or administrative, interface.

1. Go to `configuration net datalinks`.
2. Select a datalink.

   Use the `show` command to determine whether you want to modify an existing class `vnic` datalink or create a new datalink for this interface.

   To create a new datalink, perform the following steps. See also Configuring a Network Datalink (CLI).

   a. Enter the `vnic` command.

   ```
   hostname:configuration net datalinks> vnic
   ```

   b. Set a name for the datalink.

   You might want the datalink name to indicate that this is for a management interface.

   ```
   hostname:configuration net datalinks device (uncommitted)> set label=this-hostname-mgmt-dl
                             label = this-hostname-mgmt-dl (uncommitted)
   ```

   c. Optional: Set the MTU.

   By default, `mtu`, or Max transmission unit (MTU), is set to `1500`.

   d. Optional: Set a VLAN ID (`id`).

   e. Commit the changes.

```
hostname:configuration net datalinks device (uncommitted)> commit
hostname:configuration net datalinks> show
Datalinks:

DATALINK       CLASS       LINKS       STATE    ID       LABEL
igb0           device      igb0        up       -        datalink1
vnic1          vnic        igb0        up       -        this-hostname-mgmt-dl
```

3. Go to `configuration net interfaces` and enter the `ip` command.

   ```
   hostname:configuration net interfaces> ip
   ```

4. Set the name of this interface.

   Set the value of the `label` property to the name of this interface. You might want to use the name of the host for this VNIC. If the `admin` property is set to `true`, you might want to indicate that this is a management interface.

   ```
   hostname:configuration net interfaces ip (uncommitted)> set label="this-hostname-
   mgmt-if"
                            label = this-hostname-mgmt-if (uncommitted)
   ```

5. Set the `enable` property.
6. Make sure the admin property is set to `true`.
7. Set a datalink for this interface.

   Set the value of the `links` property to the name of an existing datalink. Use the `available` command to see a list of existing datalinks.

   ```
   hostname:configuration net interfaces ip (uncommitted)> set links=vnic1
                             links = vnic1 (uncommitted)
   ```

8. Specify the addresses for this interface.

   Set the `v4addrs` property or the `v6addrs` property to appropriate addresses for this interface. Enter one or more IP addresses in CIDR notation.

9. Commit the changes.

   ```
   hostname:configuration net interfaces ip (uncommitted)> commit
   hostname:configuration net interfaces> show
   Interfaces:

   INTERFACE    STATE    CLASS LINKS      ADDRS                  LABEL
   vnic1        up       ip    vnic1      nn.nn.nnn.nn/nn
                     this-hostname-mgmt-if
   ```

10. Enter `done`.
11. Configure routing.
    a. Go to `configuration net routing` and enter the `create` command.
    b. Set the family property to `IPv4` or `IPv6`.
    c. Set the destination IP address and the mask.
    d. Set the gateway IP address.
    e. Set the interface to the name of the interface that you created earlier in this procedure.
    f. Enter `commit`.
    g. Enter `done`.
12. If this is a clustered configuration, lock the interface.

See Locking a Cluster Management Interface (CLI). This step is optional but strongly recommended.

Next Steps

If this is a clustered configuration, repeat this procedure on the second controller. Ownership of an interface is defined by the node that the interface was created on. If an interface is needed for the second controller in a cluster, create the interface on that second controller.

Related Topics

- Configuring a Network Datalink (CLI)
- Locking a Cluster Management Interface (CLI)

Locking a Cluster Management Interface (BUI)

After initial configuration, clustered controllers are in an active-active state. When a failover occurs, the active controller takes over all non-private interfaces, and the peer controller becomes passive and inaccessible by its BUI and CLI. To maintain access to a controller regardless of its state, lock its management interface to make the management interface private. Use this procedure to lock the management interface of each controller in a cluster.

> **Note:**
>
> Failure to configure locked management interfaces on clustered controllers might result in longer fault diagnosis and resolution times.

1. From the **Configuration** menu, select **Cluster**.
2. Choose the management interface from the **Resource** list.
3. Click the padlock icon to lock the management interface to this controller.

   The interface displays a locked icon 🔒.
4. Click **APPLY** in the upper-right corner of the page.

Next Steps

Repeat this procedure on the second controller.

Related Topics

Configuring a Management Interface (BUI)

Locking a Cluster Management Interface (CLI)

After initial configuration, clustered controllers are in an active-active state. When a failover occurs, an active controller takes over all non-private interfaces, and the peer controller becomes passive and inaccessible by its BUI and CLI. To maintain access to a controller regardless of its state, lock its management interface to make it private. The following procedure locks the management interface on each clustered controller.

> **Caution:**
>
> Failure to configure locked management interfaces on clustered controllers may lead to longer than necessary fault diagnosis and resolution times.

1. Go to configuration cluster resources and enter the show command.
2. Select the management interface, prefaced with net/.
3. Lock the interface by setting the type to private.

```
hostname:configuration cluster resources net/vnic1> set type=private
                         type = private (uncommitted)
hostname:configuration cluster resources net/vnic1> commit
```

Next Steps

Repeat this procedure on the second controller.

Related Topics

Configuring a Management Interface (CLI)

Configuring Network IP Multipathing (IPMP)

Use network IP multipathing groups to provide IP addresses that will remain available in the event of an IP interface failure (such as a physical wire disconnection or a failure of the connection between a network device and its switch) or in the event of a path failure between the system and its network gateways. The system detects failures by monitoring the IP interface's underlying datalink for link-up and link-down notifications, and optionally by probing using test addresses that can be assigned to each IP interface in the group, described below. Any number of IP interfaces can be placed into an IPMP group so long as they are all on the same link (LAN, IB partition, or VLAN), and any number of highly-available addresses can be assigned to an IPMP group.

Each IP interface in an IPMP group is designated either *active* or *standby*:

- **Active** - The IP interface will be used to send and receive data so long as IPMP has determined it is functioning correctly.
- **Standby** - The IP interface will only be used to send and receive data if an active interface (or a previously activated standby) stops functioning.

Multiple active and standby IP interfaces can be configured, but each IPMP group must be configured with at least one active IP interface. IPMP will strive to activate as many standbys as necessary to preserve the configured number of active interfaces. For example, if an IPMP group is configured with two active interfaces and two standby interfaces and all interfaces are functioning correctly, only the two active interfaces will be used to send and receive data. If an active interface fails, one of the standby interfaces will be activated. If the other active interface fails (or the activated standby fails), the second standby interface will be activated. If the active interfaces are subsequently repaired, the standby interfaces will again be deactivated.

IP interface failures can be discovered by either link-based detection or probe-based detection (that is, a test address is configured).

If probe-based failure detection is enabled on an IP interface, the system will determine which target systems to probe dynamically. First, the routing table will be scanned for gateways (routers) on the same subnet as the IP interface's test address and up to five will be selected. If no gateways on the same subnet were found, the system will send a multicast ICMP probe (to 224.0.01. for IPv4 or ff02::1 for IPv6) and select the first five systems on the same subnet that respond. Therefore, for network failure detection and repair using IPMP, you should be sure that at least one neighbor on each link or the default gateway responds to ICMP echo requests. IPMP works with both IPv4 and IPv6 address configurations. In the case of IPv6, the interface's link-local address is used as the test address.

> **Note:**
>
> Do not use probe-based failure detection when there no systems (other than the cluster peer) on the same subnet as the IPMP test addresses that are configured to answer ICMP echo requests.

The system will probe selected target systems in round-robin fashion. If five consecutive probes are unanswered, the IP interface will be considered failed. Conversely, if ten consecutive probes are answered, the system will consider a previously failed IP interface as repaired. You can set the system's IPMP probe failure detection time from the IPMP screen. This time indirectly controls the probing rate and the repair interval—for instance, a failure detection time of 10 seconds means that the system will send probes at roughly two second intervals and that the system will need 20 seconds to detect a probe-based interface repair. You cannot directly control the system's selected targeted systems, though it can be indirectly controlled through the routing table.

The system will monitor the routing table and automatically adjust its selected target systems as necessary. For instance, if the system using multicast-discovered targets but a route is subsequently added that has a gateway on the same subnet as the IP interface's test address, the system will automatically switch to probing the gateway. Similarly, if multicast-discovered targets are being probed, the system will periodically refresh its set of chosen targets (for example, because some previously selected targets have become unresponsive).

For step-by-step instructions on building IPMP groups, see IPMP Configuration.

For information about private local interfaces, see Appliance Cluster Configuration.

Creating an IPMP Group Using Probe-Based and Link-State Failure Detection (BUI)

Create one or more "underlying" IP interfaces that will be used as components of the IPMP group. Each interface must have an IP address to be used as the probe source (see Creating a Single Port Interface (BUI)).

Do not use probe-based failure detection when there no systems (other than the cluster peer) on the same subnet as the IPMP test addresses that are configured to answer ICMP echo requests.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Click the **Interface** add item icon ⊕ .
3. Optional: Change the name of the interface.
4. Click the **IP MultiPathing Group** check box.
5. Click **Use IPv4 Protocol** and/or **Use IPv6 Protocol**, and specify the IP addresses for the IPMP interface.
6. Choose the interfaces created in the first step from the **Interfaces** list.
7. Set each chosen interface to be either **Active** or **Standby**.
8. Click **APPLY**.

Creating an IPMP Group Using Link-State Only Failure Detection (BUI)

Create one or more "underlying" IP interfaces with the IP address 0.0.0.0/8 to be used as the components of the IPMP group (see Creating a Single Port Interface (BUI)).

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Click the **Interface** add item icon ⊕ .
3. Optional: Change the name of the interface.
4. Click the **IP MultiPathing Group** check box.
5. Click **Use IPv4 Protocol** or/and **Use IPv6 Protocol** and specify the IP addresses for the IPMP interface.
6. Choose the interfaces created in the first step from the **Interfaces** list.
7. Set each chosen interface to be either **Active** or **Standby**.
8. Click **APPLY**.

Extending an IPMP Group (BUI)

Use this procedure to extend an IPMP group.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Hover over an interface in the **Interfaces** list.
3. Click the **move** icon ✥ and then drag and drop the device onto an IPMP interface.
4. Click **APPLY** at the top of the page to commit this configuration.

Creating an LACP Aggregated Link Interface (BUI)

Use this procedure to create an LACP aggregated link interface.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Click the **Datalinks** add item icon ⊕ .
3. Optionally set the datalink name.
4. Select **LACP Aggregation**.
5. Select two or more devices from the **Devices** list, and click **APPLY**.
6. Click the **Interfaces** add item icon ⊕ .
7. Set desired properties, choose the aggregated link from the **Datalinks** list, and click **APPLY**.
8. Click **APPLY** at the top to commit the configuration.

Extending an LACP Aggregation (BUI)

Use this procedure to extend an LACP aggregation.

1. From the **Configuration** menu, select **Network**, then **Configuration**.
2. Hover over a device in the **Devices** list.
3. Click the move icon ✥ and then drag and drop the device onto an aggregation datalink.
4. Click **APPLY** at the top of the page to commit this configuration.

Configuring Network Performance and Availability

IPMP and link aggregation are different technologies available in the appliance to achieve improved network performance as well as maintain network availability. In general, you deploy link aggregation to obtain better network performance, while you use IPMP to ensure high availability. The two technologies complement each other and can be deployed together to provide the combined benefits of network performance and availability.

In link aggregations, incoming traffic is spread over the multiple links that comprise the aggregation. Thus, networking performance is enhanced as more NICs are installed to add links to the aggregation. IPMP's traffic uses the IPMP interface's data addresses as they are bound to the available active interfaces. If, for example, all the data traffic is flowing between only two IP addresses but not necessarily over the same connection, then adding more NICs will not improve performance with IPMP because only two IP addresses remain usable.

Performance can be affected by the number of VNICs/VLANs configured on a datalink for a given device, as well as by using a VLAN ID. Configuring multiple VNICs over a given device may impact the performance of all datalinks over that device by up to five percent, even when VNICs are not in use. If more than eight VNICs/VLANs are configured over a given datalink, performance may degrade significantly. Also, if a datalink uses a VLAN ID, all datalink performance for that device may be impacted by an additional five percent.

Configuring Network Routing

The system provides a single IP routing table. When an IP packet is sent to a given destination, the system selects the route whose destination most closely matches the packet's destination address, subject to the system's multihoming policy. See Multihoming Policy.

The system uses the information in the routing entry to determine which IP interface to send the packet on and, if the destination is not directly reachable, the next-hop gateway to use.

If no routing entries match the destination, the packet is dropped.

If multiple routing entries are equally close matches, and are not otherwise prioritized by multihoming policy, the system will load-spread across those matching entries on a per-connection basis.

The system does not act as a router.

The routing table has the following properties for each route.

Table 2-8 Route Properties

BUI Property	CLI Property	Description
Destination	`destination` and `mask`	Range of IP destination addresses (in CIDR notation) that can match the route. A routing entry with a Destination value of `0.0.0.0/0` matches any packet (if no other route matches more precisely), and is known as a Default route. In the Insert Static Route dialog in the BUI, you can select Default or Network as the value of Kind of route.
Gateway	`gateway`	Next hop (IP address) to send the packet to, except for routes that are type system.
Family	`family`	Internet protocol: IPv4 or IPv6
Type	`type`	Origin of the route: static, dynamic, dhcp, system, or direct.
Status	`status`	Route status: active or inactive. Inactive status applies to a static or direct route associated with a disabled or offline IP interface.

Table 2-8 (Cont.) Route Properties

BUI Property	CLI Property	Description
Interface	`interface`	IP interface the packet will be sent on. If an IPMP interface is specified, then one of the active IP interfaces in the IPMP group will be chosen randomly on a per-connection basis and automatically refreshed if the chosen IP interface subsequently becomes unusable. Conversely, if a given IP interface is part of an IPMP group, it cannot be specified in the Interface property because such a route would not be highly available.

Routing entries come from a number of different origins, as identified by the **Type** property. Although the origin of a routing entry has no bearing on how it is used by the system, its origin does control whether and how it can be edited or deleted. The system supports the following types of routes.

Table 2-9 Supported Route Types

Type	Description
Static	Created and managed by the appliance administrator. See Adding a Static Route - BUI, CLI.
Dynamic	Created automatically by the appliance via the RIP and RIPng dynamic routing protocols (if enabled).
DHCP	Created automatically by the appliance part of enabling an IP interface that is configured to use DHCP. A DHCP route will be created for each default route provided by the DHCP server.
System	Created automatically by the appliance as part of enabling an IP interface. A system route will be created for each IP subnet the appliance can directly reach. Since these routes are directly reachable, the "gateway" field instead identifies the appliance's IP address on that subnet.
Direct	Created and managed as a network interface property: Directly Reachable Network(s). Directly reachable subnet that the local IP address is not a member of, but to which the datalink of its interface is physically connected. This improves scalability by conserving IP addresses, and could ease traffic congestion through core switches and routers. Direct routes are configured as network interfaces by using the **Configuration: Network: Configuration** BUI screen or the `configuration net interfaces` CLI context.

Static Routes

To ensure the appropriate network interfaces are used for the replication connections between source and target appliances, configure static /32 (host-specific) routes.

If you are setting up replication for a cluster configuration, select a singleton (unlocked) network interface so that following a cluster takeover or failback, the interface will move to the node where the replication work is being done. The two source cluster nodes can replicate to the same target node only if the target node provides two IP addresses, one for use by each node in the source cluster. Replicating to the same target IP address from both nodes of a source cluster is not supported.

Note:

When an interface is deleted, all routes associated with the interface are also removed.

Related Topics

- Remote Replication Workflow
- Remote Replication Concepts

Multihoming Policy

If a system is configured with more than one IP interface, then the system might have multiple equivalent routes to a given destination, forcing the system to choose which IP interface to send a packet on. Similarly, a packet may arrive on one IP interface, but be destined to an IP address that is hosted on another IP interface. The system's behavior in such situations is determined by the selected multihoming policy.

The multihoming policy value controls the system policy for accepting and transmitting IP packets when multiple IP interfaces are simultaneously enabled. The value of the multihoming policy can be loose (default), adaptive, or strict.

Table 2-10 Multihoming Policies

Policy	Description
`Loose`	Do not enforce any binding between an IP packet and the IP interface used to send or receive the packet. • An IP packet will be accepted on an IP interface as long as its destination IP address is up on the appliance. • An IP packet will be transmitted over the IP interface tied to the route that most specifically matches an IP packet's destination address, without any regard for the IP addresses hosted on that IP interface. If no eligible routes exist, drop the packet.
`Adaptive`	Identical to `Loose`, except prefer routes with a gateway address on the same subnet as the packet's source IP address. • An IP packet will be accepted on an IP interface as long as its destination IP address is up on the appliance. • An IP packet will be transmitted over the IP interface tied to the route that most specifically matches an IP packet's destination address. If multiple routes are equally specific, prefer routes that have a gateway address on the same subnet as the packet's source address. If no eligible routes exist, drop the packet.
`Strict`	Require a strict binding between an IP packet and the IP interface used to send or receive it. • An IP packet will be accepted on an IP interface as long as its destination IP address is up on that IP interface. • An IP packet will only be transmitted over an IP interface if its source IP address is up on that IP interface. To enforce this, when matching against the available routes, the appliance will ignore any routes that have gateway addresses on a different subnet from the packet's source address. If no eligible routes remain, drop the packet.

When selecting the multihoming policy, a key consideration is whether any of the appliance's IP interfaces will be dedicated to administration (for example, for dedicated BUI access) and thus accessed over a separate administration network. In particular, if a default route is created to provide remote access to the administration network, and a separate default route is created to provide remote access to storage protocols, then the default system policy of `Loose` might cause the administrative default route to be used for storage traffic. By switching the policy to `Adaptive` or `Strict`, the appliance will consider the IP address associated with the request as part of selecting the route for the reply. If no route can be found on the same IP interface, the `Adaptive` policy will cause the system to use any available route, whereas the `Strict` policy will cause the system to drop the packet.

Adding a Static Route (BUI)

After defining the static route from the source appliance to the target appliance, repeat these steps on the target appliance to define the static route from the target back to the source.

1. Source appliance: From the **Configuration** menu, select **Network**, then **Routing**.
2. Click the add item icon ⊕ .
3. In the **Insert Static Route** dialog box, set properties.
   - Select a **Family protocol** (**IPv4** or **IPv6**).
   - Select **Network for Kind**.
   - Specify a **Destination address** (the IP address of the target appliance) and a **Gateway address**.
   - Select an **Interface**.
4. Click **ADD**.

   The new route will appear in the table.
5. Target appliance: From the **Configuration** menu, select **Network**, then **Routing**.
6. Repeat step 2 through step 4.

   The value of **Destination** is the IP address of the source appliance.

Adding a Static Route (CLI)

After defining the static route from the source appliance to the target appliance, repeat these steps on the target appliance to define the static route from the target back to the source.

Use a static `/32` (host-specific) route to the target system IP address via the dedicated network interface.

1. Source appliance: Go to `configuration services routing`.
2. Enter `create`.
3. Type `show` or `get` to list properties.
4. Set the properties.
   - Set the `family`, `destination`, `mask`, `gateway`, and `interface`.
   - The value of `destination` is the IP address of the target appliance.
   - The value of `mask` should be `32` to indicate that this is a static, host-specific route.
   - You can use tab completion to see a list of values for `family` and `interface`.

5. Enter `commit`.
6. Target appliance: Go to `configuration services routing`.
7. Repeat step 2 through step 5.

   The value of `destination` is the IP address of the source appliance.

To ensure traffic is routed through the correct source and target interfaces, use the `traceroute` command.

Deleting a Static Route (BUI)

Use this procedure to delete a static route.

1. From the **Configuration** menu, select **Network**, then **Routing**.
2. Hover over the route entry, then click the trash icon 🗑 .

Deleting a Static Route (CLI)

Use this procedure to delete a static route.

1. Go to `configuration net routing`.
2. Type `show` or `get` to list routes and route names.
3. Enter `destroy` *`route-name`*.

Changing the Multihoming Property to Strict (BUI)

Use this procedure to change the `multihoming` property to `Strict`.

1. From the **Configuration** menu, select **Network**, then **Routing**.
2. In the **Multihoming model** area at the top of the page, click **Strict**.

Changing the Multihoming Property to Strict (CLI)

Use this procedure to change the `multihoming` property to `strict`.

1. Go to `configuration net routing`.
2. Enter `set multihoming=strict`.
3. Enter `commit`.

Configuring Storage

Oracle ZFS Storage Appliance uses storage pools to manage physical storage devices. After configuring these pools based on physical characteristics and the desired level of data redundancy, you can store filesystems and LUNs, collectively known as shares, in these pools. Shares, which are contained in projects, automatically grow within the disk space allocated to the pool, and pools can span multiple storage devices. Although there is no need to statically size shares, you can control space usage using quotas and reservations. For more information, see Space Management for Shares.

To configure and manage storage, use these tasks:

- Creating a Storage Pool - BUI, CLI
- Importing an Existing Storage Pool - BUI, CLI

- Configuring an All-Flash Storage Pool - BUI, CLI
- Adding a Disk Shelf to an Existing Storage Pool - BUI, CLI
- Adding a Cache, Meta, or Log Device to an Existing Storage Pool - BUI, CLI
- Removing a Cache or Log Device from an Existing Storage Pool - BUI, CLI
- Unconfiguring a Storage Pool - BUI, CLI
- Renaming a Storage Pool - BUI, CLI
- Scrubbing a Storage Pool - Manual - BUI, CLI
- Scrubbing a Storage Pool - Scheduled - BUI, CLI
- Viewing Pool and Device Status - BUI, CLI

To understand storage basics, use these topics:

- Storage Pool Concepts
- Data Profiles for Storage Pools
- Space Management for Shares

Creating a Storage Pool (BUI)

Storage pools store data and can be created during or after initial configuration. Pools can contain data drives and log, read cache, and meta devices.

Pools can be encrypted. Use this procedure to create an unencrypted pool. To create an encrypted pool, see Creating an Encrypted Pool (BUI). You cannot add encryption information to a pool that was already created as unencrypted.

This procedure assumes that initial configuration has been completed. Creating and configuring a storage pool is a two-step process. First, the storage devices are verified for presence and minimum functionality, and you assign drives or even entire disk shelves to the pool. Second, you select a profile for the drives based on your storage needs. If for some reason a pool is unconfigured, you can import it as described in Importing an Existing Storage Pool (BUI).

> **Note:**
>
> After a storage pool is created and a file is actively retained within that pool using the mandatory file retention policy, the pool cannot be unconfigured until all mandatory file retention has expired. For more information, see File Retention Management.

To reduce redundant data, which can be especially prevalent in replication workloads, consider the benefits of using deduplication. Allocate meta devices if deduplication will be enabled for projects or shares in this pool. For more information, see Data Deduplication. There is also an all-flash storage pool, which utilizes SSDs as data devices and optional log devices, but does not contain read cache or meta devices. See Configuring an All-Flash Storage Pool (BUI).

Before You Begin

- For recommendations on how many drives to select per pool, see Number of Devices per Pool.
- To understand the different data profiles, see Data Profiles for Storage Pools.

- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** menu, select **System**.
- To use the enhanced data deduplication feature in a storage pool, upgrade to software release OS8.7.0 or later and accept all deferred updates, including Data Deduplication v2. See Data Deduplication v2 Deferred Update in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

1. From the **Configuration** menu, select **Storage**.
2. Next to **Available Pools**, click the add icon ⊕ .
3. Type a name for the storage pool and click **APPLY**.
4. Select the number of data drives for the storage pool for each disk shelf. You can also select available log, cache, and meta devices.

   For more information on log, cache, and meta devices, see Data Profiles for Storage Pools.

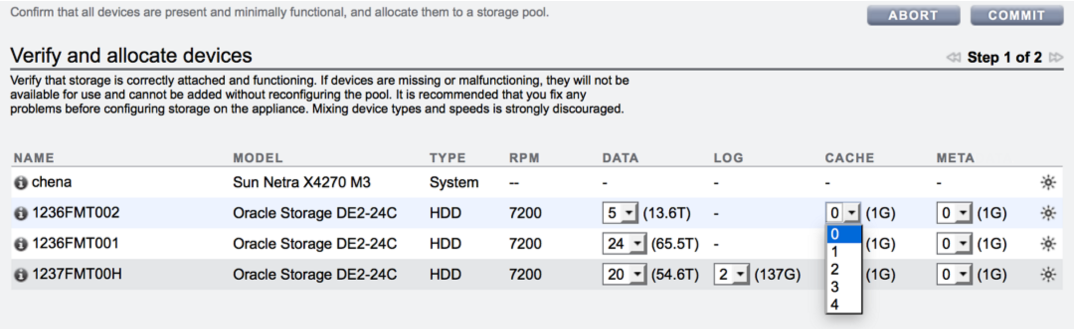

Confirm that all devices are present and minimally functional, and allocate them to a storage pool.

Verify and allocate devices

Verify that storage is correctly attached and functioning. If devices are missing or malfunctioning, they will not be available for use and cannot be added without reconfiguring the pool. It is recommended that you fix any problems before configuring storage on the appliance. Mixing device types and speeds is strongly discouraged.

NAME	MODEL	TYPE	RPM	DATA	LOG	CACHE	META
chena	Sun Netra X4270 M3	System	--	-	-	-	-
1236FMT002	Oracle Storage DE2-24C	HDD	7200	5 (13.6T)	-	0 (1G)	0 (1G)
1236FMT001	Oracle Storage DE2-24C	HDD	7200	24 (65.5T)	-	(1G)	0 (1G)
1237FMT00H	Oracle Storage DE2-24C	HDD	7200	20 (54.6T)	2 (137G)	(1G)	0 (1G)

Caution:

Once a data disk has been added to a pool, it cannot be removed without destroying the pool entirely and losing all data.

If all attached disk shelves do not appear, click **ABORT**, check the disk shelf cabling and power, and begin this procedure again.

- If all drives are the same size or rotational speed, or if one size is selected among multiple sizes, the maximum number of drives available is allocated by default. If the storage device contains drives of different rotational speeds or models, no drives are allocated by default.
- It is strongly recommended that pools include only devices of the same size and rotational speed to provide consistent performance characteristics.
- Monitor or limit space usage because you may experience reduced performance when pools approach full capacity.

5. Click **COMMIT**.

   The drives are allocated to the storage pool, and verified for presence and minimum functionality. If verification fails, click **ABORT**, fix the problem, and begin this procedure again. If you allocate a pool with missing or failed devices, you will not be able to add the missing or failed devices later.

6. On the **Choose Storage Profile** screen, select the data profile that meets your reliability, availability, serviceability, and performance goals.

   For a description of each profile, click on the data profile name, or see Data Profiles for Storage Pools.

   > **Note:**
   >
   > For mandatory file retention, the storage pool profile must provide redundancy. Therefore, the striped profile cannot be used with storage pools that will contain files with the mandatory file retention policy. For more information, see File Retention Management.

7. If you allocated log, cache, or meta devices, select the appropriate profiles.

   - For log devices, click **Log Profile** and select either the **mirrored** or **striped** profile. If you allocated an even number of log devices to the pool, select the **mirrored** profile.

     > **Caution:**
     >
     > A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see Data Profiles for Storage Pools.

   - For cache devices, the profile is always **striped**, as shown under **Cache Profile**.
   - For meta devices, click **Metadata Profile** and select either the **mirrored** or **striped** profile.

   > **Note:**
   >
   > Once meta devices are added to a storage pool, they cannot be removed from the pool.

8. Click **COMMIT**.

Related Topics

- Data Profiles for Storage Pools
- Importing an Existing Storage Pool (BUI)
- Adding a Disk Shelf to an Existing Storage Pool (BUI)
- Renaming a Storage Pool (BUI)
- Storage Pool Concepts
- Data Deduplication

Creating a Storage Pool (CLI)

Storage pools store data and can be created during or after initial configuration. Pools can contain data drives and log, read cache, and meta devices.

Pools can be encrypted. Use this procedure to create an unencrypted pool. To create an encrypted pool, see Creating an Encrypted Pool (CLI). You cannot add encryption information to a pool that was already created as unencrypted.

This procedure assumes that initial configuration has been completed. Creating and configuring a storage pool is a two-step process. First, the storage devices are verified for presence and minimum functionality, and you assign drives or even entire disk shelves to the pool. Second, you select a profile for the drives based on your storage needs. If for some reason a pool is unconfigured, you can import it as described in Importing an Existing Storage Pool (CLI).

> **Note:**
>
> After a storage pool is created and a file is actively retained within that pool using the mandatory file retention policy, the pool cannot be unconfigured until all mandatory file retention has expired. For more information, see File Retention Management.

To reduce redundant data, which can be especially prevalent in replication workloads, consider the benefits of using deduplication. Allocate meta devices if deduplication will be enabled for projects or shares in this pool. For more information, see Data Deduplication. There is also an all-flash storage pool, which utilizes SSDs as data devices and optional log devices, but does not contain read cache or meta devices. See Configuring an All-Flash Storage Pool (CLI).

Before You Begin

- For recommendations on how many drives to select per pool, see Number of Devices per Pool.
- To understand the different data profiles, see Data Profiles for Storage Pools.
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.
- To use the enhanced data deduplication feature in a storage pool, upgrade to software release OS8.7.0 or later and accept all deferred updates, including Data Deduplication v2. See Data Deduplication v2 Deferred Update in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

1. Go to `configuration storage`.
2. Enter `config` and a name for the new storage pool.

   ```
   hostname: configuration storage> config pool0
   hostname: configuration storage (pool0) verify>
   ```

3. Enter `show` to see the device information for the pool:

   ```
   hostname:configuration storage (pool0) verify> show
   ID STATUS  ALLOCATION  DATA   LOG      CACHE   META     RPM
   0      ok     custom      0   0          0/4    0/4    1.86T
   1      ok     custom      0   0/2  34G     0      0    15000
   2      ok     custom      0   0/2  34G     0      0    15000
   ```

4. Enter `set` and the disk shelf or controller ID, and the number of data drives to use. You can also select available cache, meta, and log devices.

   For more information on log, cache, and meta devices, see Data Profiles for Storage Pools.

> **Caution:**
>
> Once a data disk has been added to a pool, it cannot be removed without destroying the pool entirely and losing all data.

ID "0" is the controller, and the remaining IDs are the disk shelves. In the following example, `1-data=8` allocates eight data drives from the first disk shelf.

```
hostname:configuration storage (pool1) verify> set 1-data=8
                            1-data = 8
```

This example allocates one cache device from the controller:

```
hostname:configuration storage (pool1) verify> set 0-cache=1
                            0-cache = 1
```

This example allocates one meta device from the controller:

```
hostname:configuration storage (pool1) verify> set 0-meta=1
                            0-meta = 1
```

5. Enter `done`.

   ```
   hostname:configuration storage (pool1) verify> done
   ```

6. Enter `show` to display the profile.

   ```
   hostname:configuration storage (pool1) config> show
   PROFILE                 CAPCITY  NSPF  DESCRIPTION
   log_profile = log_stripe   17G   no    Striped log
   ```

   If you allocated cache devices to the pool, the profile is always `striped`.

   > **Note:**
   >
   > For mandatory file retention, the storage pool profile must provide redundancy. Therefore, the striped profile cannot be used with storage pools that will contain files with the mandatory file retention policy. For more information, see File Retention Management.

7. If you allocated log devices to the pool, enter `set log_profile=` and set the log profile to either `log_mirror` or `log_stripe`. Use `log_mirror` if the pool contains an even number of log devices.

   > **Caution:**
   >
   > A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see Data Profiles for Storage Pools.

   ```
   hostname:configuration storage (pool1)> set log_profile=log_mirror
   ```

8. If you allocated meta devices to the pool, enter `set meta_profile=` and set the meta profile to either `meta_mirror` or `meta_stripe`.

```
hostname:configuration storage (pool1)> set meta_profile=meta_mirror
```

9. Enter `done` to complete the task.

```
hostname:configuration storage (pool1)> done
```

Related Topics

- Data Profiles for Storage Pools
- Importing an Existing Storage Pool (CLI)
- Adding a Disk Shelf to an Existing Storage Pool (CLI)
- Renaming a Storage Pool (CLI)
- Storage Pool Concepts
- Data Deduplication

Importing an Existing Storage Pool (BUI)

The import action allows you to import an unconfigured storage pool. A storage pool can be unconfigured because of an inadvertent action, factory reset, or service operation to recover user data. Importing a storage pool requires scanning all attached storage devices and discovering any existing state. This can take a significant amount of time, during which no other storage configuration activities can take place.

Before You Begin

Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check whether an upgrade is in progress, from the **Maintenance** menu, select **System**.

1. From the **Configuration** menu, select **Storage**.

   A list of storage pools is displayed, including some identifying characteristics. If the storage has been destroyed or is incomplete, the storage pool is not importable. Unlike storage configuration, the storage pool name is not initially shown, but it is shown after selecting the storage pool.

2. Click **IMPORT**.
3. Select the storage pool you want to import.

   By default, the previous storage pool names are displayed.

4. To rename the storage pool, click the pool name and change it.
5. Click **COMMIT**.

Related Topics

- Unconfiguring a Storage Pool (BUI)
- Renaming a Storage Pool (BUI)

Importing an Existing Storage Pool (CLI)

The import action allows you to import an unconfigured storage pool. A storage pool can be unconfigured because of an inadvertent action, factory reset, or service operation to recover user data. Importing a storage pool requires iterating over all attached storage devices and discovering any existing state. This can take a significant amount of time, during which no other storage configuration activities can take place.

Before You Begin

Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check whether an upgrade is in progress, navigate to `maintenance system updates`.

1. Go to `configuration storage`.
2. Enter `import`.

   ```
   hostname:configuration storage (pool0)> import

   Search for storage. Begin the process of searching for existing storage pools.

   Subcommands that are valid in this context:

      help [topic]           => Get context-sensitive help. If [topic] is specified,
                                it must be one of "builtins", "commands", "general",
                                "help" or "script".

      show                   => Show information pertinent to the current context

      abort                  => Abort this task (potentially resulting in a
                                misconfigured system)

      done                   => Finish operating on "discover"

   hostname:configuration storage (pool0) discover>
   ```

3. Enter `done`.
4. Enter `show`.

   ```
   hostname:configuration storage (pool0)> show
   Pools:

      POOL             OWNER         DATA PROFILE  LOG PROFILE  STATUS   ERRORS
   -> pool0          hostname        mirror        log_stripe   online   0
      pool1          hostname        -             -            exported -

   Properties:
                              pool = pool0
                            status = online
                            errors = 0
                             owner = hostname
                           profile = mirror
                       log_profile = log_stripe
                     cache_profile = cache_stripe
                             scrub = none requested
   ```

5. Enter `set pool=` and the name of the pool you want to import.

   > **Note:**
   >
   > If you have a single pool, the pool name is not displayed, but it is selected.

   ```
   hostname:configuration storage select> set pool=pool1
                               pool = pool1
   ```

   A message reminds you to verify that storage is correctly attached and functioning.

6. Enter `done`.

Related Topics

- Unconfiguring a Storage Pool (CLI)
- Renaming a Storage Pool (CLI)

Configuring an All-Flash Storage Pool (BUI)

An all-flash storage pool utilizes SSDs as data devices and optional log devices, but does not contain read cache or meta devices. All-flash pools are suitable for virtualization environments or backup workloads.

Before You Begin

- Follow the cabling guidelines for all-flash shelves described in Cabinet and Cabling Guidelines in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.
- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see Storage Pool Concepts.
- To understand the different data profiles, see Data Profiles for Storage Pools.

> **Note:**
>
> Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** menu, select **System**.

1. From the **Configuration** menu, select **Storage**.
2. Click the add icon ⊕ above the list of storage pools.
3. From the **Verify and allocate devices** screen, select storage type `SSD`, and then select a device size.

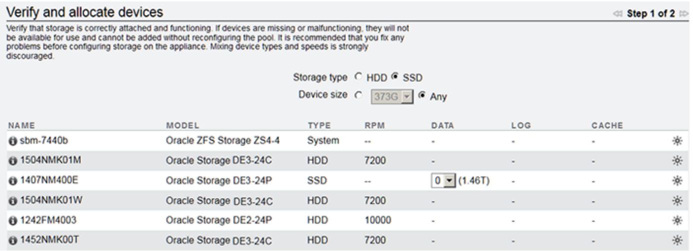

Verify and allocate devices

Verify that storage is correctly attached and functioning. If devices are missing or malfunctioning, they will not be available for use and cannot be added without reconfiguring the pool. It is recommended that you fix any problems before configuring storage on the appliance. Mixing device types and speeds is strongly discouraged.

Storage type ○ HDD ● SSD

Device size ○ 373G ● Any

NAME	MODEL	TYPE	RPM	DATA	LOG	CACHE
sbm-7440b	Oracle ZFS Storage ZS4-4	System	--	-	-	-
1504NMK01M	Oracle Storage DE3-24C	HDD	7200	-	-	-
1407NM400E	Oracle Storage DE3-24P	SSD	--	0 (1.46T)	-	-
1504NMK01W	Oracle Storage DE3-24C	HDD	7200	-	-	-
1242FM4003	Oracle Storage DE2-24P	HDD	10000	-	-	-
1452NMK00T	Oracle Storage DE3-24C	HDD	7200	-	-	-

4. For each SSD disk shelf, select the number of drives to include in the pool.

   > **Note:**
   >
   > An all-flash pool cannot contain read cache devices or meta devices.

5. Optional: Select log devices to add to the all-flash pool.
6. Click **COMMIT**.
7. On the **Configure Added Storage** screen, select the data profile appropriate for your workload that balances performance, availability, and capacity.

   For a description of available profiles, see Data Profiles for Storage Pools.

> **Note:**
>
> For mandatory file retention, the storage pool profile must provide redundancy. Therefore, the striped profile cannot be used with storage pools that will contain files with the mandatory file retention policy. For more information, see File Retention Management.

8. Optional: If you allocated log devices, select an appropriate profile.
9. Click **COMMIT**.

Related Topics

- All-Flash Storage Configuration
- Setting a Threshold Alert for SSD Endurance (BUI) in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*

Configuring an All-Flash Storage Pool (CLI)

An all-flash storage pool utilizes SSDs as data devices and optional log devices, but does not contain read cache or meta devices. All-flash pools are suitable for virtualization environments or backup workloads.

Before You Begin

- Follow the cabling guidelines for all-flash shelves described in Cabinet and Cabling Guidelines in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.
- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see Storage Pool Concepts.
- To understand the different data profiles, see Data Profiles for Storage Pools.

> **Note:**
>
> Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.

1. Verify SSDs are correctly attached and functioning.

   If any devices are missing or malfunctioning, make the necessary corrections.

   > **Note:**
   >
   > An all-flash pool cannot contain read cache devices or meta devices.

2. Go to `configuration storage`, enter `config` and a unique name for the pool:

   ```
   hostname:configuration storage>
                config allflashpool
   ```

   Instructions and subcommands that can be used in this context are displayed.

3. Enter `show` to view the available devices for the pool.

```
hostname:configuration storage verify>
              show

       ID    STATUS   ALLOCATION    DATA          LOG           CACHE       RPM
TYPE
       --    -------   ----------   -----------   -----------   ----------- -----
------
        0    ok       custom        0             0             0
system
        1    ok       custom        0/7   3.46T   0/2   373G    0
ssd
        2    ok       custom        0/24  6.55T   0             0
ssd
                                    -----------   -----------   -----------
                                    0             0             0
```

4. Enter `help properties` to list the available properties:

```
hostname:configuration storage verify>
              help properties

  0                    => Chassis 0
  1-data               => Chassis 1 data
  1-log                => Chassis 1 log
  2                    => Chassis 2
  2-data               => Chassis 2 data
```

5. Enter `set` [`1-data=` | `2-data=`] to assign the devices to a pool, as shown in this example:

```
hostname:configuration storage verify>
              set 1-data=3 2-data=3

                        1-data = 3
                        2-data = 3
```

   This example assigns 3 devices from chassis 1 (`1-data=3`) and 3 devices from chassis 2 (`2-data=3`) to the pool.

6. Optional: Select log devices to add to the all-flash pool.

7. Enter `done` to close verify.

```
hostname:configuration storage verify>
              done
```

8. Enter `show` to view the available storage profile types:

```
hostname:configuration storage config>
              show
```

9. Enter `set profile=` to specify the data profile appropriate for your workload, that balances performance, availability, and capacity.

   For a description of available profiles, see Data Profiles for Storage Pools.

```
hostname:configuration storage config>
              set profile=
```

> **Note:**
>
> For mandatory file retention, the storage pool profile must provide redundancy. Therefore, the striped profile cannot be used with storage pools that will contain files with the mandatory file retention policy. For more information, see File Retention Management.

10. Optional: If you allocated log devices, select an appropriate profile.
11. Enter `done`.

```
hostname:configuration storage config>
                 done
```

Related Topics

- All-Flash Storage Configuration
- Setting a Threshold Alert for SSD Endurance (CLI) in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*

Adding a Disk Shelf to an Existing Storage Pool (BUI)

Use the following task to add a disk shelf to an existing storage pool.

Before You Begin

- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see Storage Pool Concepts.
- You must select the same data profile currently used in the existing pool. To understand the different data profiles, see Data Profiles for Storage Pools.
- If there is insufficient storage to configure the system for the data profile and its options, some attributes may not be supported. For example, it is impossible to preserve NSPF characteristics when adding a single disk shelf to a double parity RAID configuration with the NSPF option. You can add the disk shelf, but cannot use the NSPF option.
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** menu, select **System**.

> **Caution:**
>
> Once a disk has been added to a pool, it cannot be removed without destroying (unconfiguring) the pool entirely and losing all data.

> **Note:**
>
> After a storage pool is created and a file is actively retained within that pool using the mandatory file retention policy, the pool cannot be unconfigured until all mandatory file retention has expired. For more information, see File Retention Management.

1. Install the new disk shelf using Adding a New Disk Shelf in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.
2. From the **Configuration** menu, select **Storage**.
3. From the **Available Pools** list, select an online pool to which to add the disk shelf.
4. Click **ADD**.
5. For this disk shelf, select the number of data drives for the storage pool.

   If the new disk shelf does not appear, click **ABORT**, check the disk shelf cabling and power, and begin this procedure again.

   - If all drives are the same size or rotational speed, or if one size is selected among multiple sizes, the maximum number of drives available is allocated by default. If the storage device contains drives of different rotational speeds or models, no drives are allocated by default.
   - It is strongly recommended that pools include only devices of the same size and rotational speed to provide consistent performance characteristics.
   - Monitor or limit space usage because you may experience reduced performance when pools approach full capacity.
6. Optional: Add any cache or log devices from the disk shelf to the pool.
7. Click **COMMIT**.
8. For data drives, select the same data profile used in the existing pool.
9. If you allocated log or cache devices, select the appropriate profiles.
   - For log devices, click **Log Profile** and select either the **mirrored** or **striped** profile. If you allocated an even number of log devices to the pool, select the **mirrored** profile.

     > **Caution:**
     >
     > A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see Data Profiles for Storage Pools.

   - For cache devices, the profile is always striped, as shown under **Cache Profile**.
10. Click **COMMIT**.

Related Topics

- Unconfiguring a Storage Pool (BUI)
- Adding a Cache, Meta, or Log Device to an Existing Storage Pool (BUI)

Adding a Disk Shelf to an Existing Storage Pool (CLI)

Use the following task to add a disk shelf to an existing storage pool.

Before You Begin

- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see Storage Pool Concepts.
- You must select the same data profile currently used in the existing pool. To understand the different data profiles, see Data Profiles for Storage Pools.

- If there is insufficient storage to configure the system for the data profile and its options, some attributes may not be supported. For example, it is impossible to preserve NSPF characteristics when adding a single disk shelf to a double parity RAID configuration with the NSPF option. You can add the disk shelf, but cannot use the NSPF option.
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.

> **Caution:**
>
> Once a disk has been added to a pool, it cannot be removed without destroying (unconfiguring) the pool entirely and losing all data.

> **Note:**
>
> After a storage pool is created and a file is actively retained within that pool using the mandatory file retention policy, the pool cannot be unconfigured until all mandatory file retention has expired. For more information, see File Retention Management.

1. Install the new disk shelf using Adding a New Disk Shelf in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.
2. Go to `configuration storage`.
3. If you have multiple pools, a default pool is selected and displayed. If this is not the pool to which you want to add the device, enter `set pool=` and specify another online pool.

   > **Note:**
   >
   > If you have a single pool, the pool name is not displayed, but it is selected.

   ```
   hostname:configuration storage (pool0)> set pool=pool1
                            pool = pool1
   ```

   A message reminds you to verify that the device is correctly installed. Note that mixing device types and speeds is strongly discouraged.
4. Enter `add`.

   ```
   hostname:configuration storage (pool1)> add
   ```
5. Enter `show` to see the device information for the pool:

   ```
   hostname:configuration storage (pool1) verify> show
   ID STATUS  ALLOCATION  DATA   LOG      CACHE   RPM
   0      ok     custom      0   0          0/4   1.86T
   1      ok     custom      0   0/2  34G      0   15000
   2      ok     custom      0   0/2  34G      0   15000
   ```
6. Specify which disk shelf or the controller and the number of data drives to use.

   ID "0" is the controller, and the remaining IDs are the disk shelves. In the following example, `1-data=8` allocates eight data drives from the first disk shelf.

```
hostname:configuration storage (pool1) verify> set 1-data=8
                            1-data = 8
```

7. Optional: Specify which disk shelf or the controller and the number of log or cache devices to use.

   ID `0` is the controller, and the remaining IDs are the disk shelves. In the following example, `set 0-cache=1` allocates one cache device from the controller:

   ```
   hostname:configuration storage (pool1) verify> set 0-cache=1
                               0-cache = 1
   ```

8. Enter `done`.

   ```
   hostname:configuration storage (pool1) verify> done
   ```

   The storage devices are verified for presence and minimum functionality. If verification fails, fix the problem, and begin this procedure again. If you allocate a pool with missing or failed devices, you will not be able to add the missing or failed devices later.

9. Enter `show` to display the profile.

   ```
   hostname:configuration storage (pool1) config> show
   PROFILE   CAPACITY   NSPF   DESCRIPTION
   log_profile   17G     no     Striped log
   ```

10. Enter the same data profile as the remainder of the pool by entering `set profile=` and the profile name.

11. Enter `done`.

12. If you allocated log devices to the pool, enter `set log_profile=`[`log_mirror | log_stripe`]. Use `log_mirror` if the pool contains an even number of log devices.

    ```
    hostname:configuration storage (pool1)> set log_profile=log_mirror
    ```

    > **Note:**
    >
    > If you allocated cache devices to the pool, the profile is always striped.

13. Enter `done`.

Related Topics

- Unconfiguring a Storage Pool (CLI)
- Adding a Cache, Meta, or Log Device to an Existing Storage Pool (CLI)

Adding a Cache, Meta, or Log Device to an Existing Storage Pool (BUI)

Use the following task to add a log, read cache, or meta device to an existing storage pool.

Before You Begin

- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see Storage Pool Concepts.
- You must select the same data profile currently used in the existing pool. To understand the different data profiles, see Data Profiles for Storage Pools.
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** menu, select **System**.

- A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 or later.

1. Install the new log, read cache, or meta device into the first available and appropriate slot. To determine the appropriate slot, see Disk Shelf Configurations in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.
2. From the **Configuration** menu, select **Storage**.
3. From the **Available Pools** list, select an online pool to which to add the device.
4. Click **ADD**.
5. Select the device to add to the pool, and click **COMMIT**.
6. Select the appropriate profiles.
   - For log devices, click **Log Profile** and select either the **mirrored** or **striped** profile. Use the **mirrored** profile if the pool now contains an even number of log devices.

     > **Caution:**
     >
     > A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see Data Profiles for Storage Pools.

   - For cache devices, the profile is always **striped**, as shown under **Cache Profile**.
   - For meta devices, click **Metadata Profile** and select either the **mirrored** or **striped** profile. The **striped** profile is recommended for better performance in the event of a meta device failure.
7. Click **COMMIT**.

Related Topics

- Removing a Cache or Log Device from an Existing Storage Pool (BUI)
- Adding a Disk Shelf to an Existing Storage Pool (BUI)
- Data Profiles for Storage Pools
- Storage Pool Concepts

Adding a Cache, Meta, or Log Device to an Existing Storage Pool (CLI)

Use the following task to add a read cache device or log device to an existing storage pool.

Before You Begin

- For recommendations on how many drives to select per pool, as well as other considerations and guidelines, see Storage Pool Concepts.
- You must select the same data profile currently used in the existing pool. To understand the different data profiles, see Data Profiles for Storage Pools.
- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.
- A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 or later.

1. Install the new log, read cache, or meta device into the first available and appropriate slot. To determine the appropriate slot, see Disk Shelf Configurations in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.
2. Go to `configuration storage`.
3. If you have multiple pools, a default pool is selected and displayed. If this is not the pool to which you want to add a device, enter `set pool=` and specify another online pool.

   > **Note:**
   >
   > If you have a single pool, the pool name is not displayed, but it is selected.

   ```
   hostname:configuration storage (pool0)> set pool=pool1
                               pool = pool1
   ```

   A message reminds you to verify that storage is correctly attached and functioning.
4. Enter `add`.

   ```
   hostname:configuration storage (pool1)> add
   ```
5. Enter `show` to display device information for the pool.

   ```
   hostname:configuration storage (pool1) verify> show
   ID STATUS  ALLOCATION  DATA   LOG      CACHE   META     RPM
   0      ok     custom      0   0          0/4    0/4    1.86T
   1      ok     custom      0   0/2  34G     0      0    15000
   2      ok     custom      0   0/2  34G   0/2      0    15000
   ```
6. Enter `set` and use tab completion to see if cache, meta, and log devices are available.

   ```
   hostname:configuration storage (pool1) verify> set
   0-cache  1-data   2-cache  2-meta  2-log
   ```
7. Enter `set` and the disk shelf or controller ID, and the number of log, cache, or meta devices to use.

   ID `0` is the controller, and the remaining IDs are the disk shelves. In the following example, `2-log=1` allocates one log device from the second disk shelf.

   ```
   hostname:configuration storage (pool1) verify> set 2-log=1
                               2-log = 1
   ```

   > **Note:**
   >
   > A value of `1-log=2` would allocate two log devices from the first disk shelf.

   This example allocates one cache device from the second disk shelf.

   ```
   hostname:configuration storage (pool1) verify> set 2-cache=1
                               2-cache = 1
   ```

   This example allocates one meta device from the second disk shelf:

   ```
   hostname:configuration storage (pool1) verify> set 2-meta=1
                               2-meta = 1
   ```
8. Enter `done`.

```
hostname:configuration storage (pool1) verify> done
```

9. Enter `show` to display the profile.

```
hostname:configuration storage (pool1) config> show
PROFILE                      CAPCTY   NSPF   DESCRIPTION
log_profile = log_stripe       17G      no   Striped log
```

> **Note:**
>
> If you allocated cache devices to the pool, the profile is always striped.

10. If you allocated log devices to the pool, enter `set log_profile=` and set the log profile to either `log_mirror` or `log_stripe`. Use `log_mirror` if the pool now contains an even number of log devices.

> **Caution:**
>
> A double failure can cause loss of data from a log in a striped configuration. It is highly recommended to configure a mirrored log profile for added redundancy. For more information, see Data Profiles for Storage Pools.

```
hostname:configuration storage (pool1)> set log_profile=log_mirror
```

11. If you allocated meta devices to the pool, enter `set meta_profile=` and set the meta profile to either `meta_mirror` or `meta_stripe`.

```
hostname:configuration storage (pool1)> set meta_profile=meta_mirror
```

12. Enter `done` to complete the task.

```
hostname:configuration storage (pool1)> done
```

Related Topics

- Removing a Cache or Log Device from an Existing Storage Pool (CLI)
- Adding a Disk Shelf to an Existing Storage Pool (CLI)
- Data Profiles for Storage Pools
- Storage Pool Concepts

Removing a Cache or Log Device from an Existing Storage Pool (BUI)

Use the following task to remove a read cache or log device from an existing storage pool. This capability is useful when preparing for a system update that requires the removal of certain cache devices.

> **Note:**
>
> Meta devices cannot be removed from a storage pool.

If a pool has cache devices on both controllers of a clustered configuration, you must perform this procedure on each controller.

To add a device to a different, existing storage pool, see Adding a Cache, Meta, or Log Device to an Existing Storage Pool (BUI).

Before You Begin

Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** menu, select **System**.

1. From the **Configuration** menu, select **Storage**.
2. From the **Available Pools** list, select an online pool from which to remove the device.
3. Click **REMOVE**.
4. Select the number of log and cache devices to be removed from the storage pool.

   > **Note:**
   >
   > If the log devices use a mirrored profile, a message reminds you to select an even number of log devices to remove. If they use a striped profile, you may remove an even or odd number of devices.

5. Click **COMMIT**.

Related Topics

Adding a Cache, Meta, or Log Device to an Existing Storage Pool (BUI)

Removing a Cache or Log Device from an Existing Storage Pool (CLI)

Use the following task to remove a read cache or log device from an existing storage pool. This capability is useful when preparing for a system update that requires the removal of certain cache devices.

> **Note:**
>
> Meta devices cannot be removed from a storage pool.

If a pool has cache devices on both controllers of a clustered configuration, you must perform this procedure on each controller.

To add a device to a different, existing storage pool, see Adding a Cache, Meta, or Log Device to an Existing Storage Pool (CLI).

Before You Begin

Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.

1. Go to `configuration storage`.
2. If you have multiple pools, a default pool is displayed and selected. If this is not the pool to which you want to add the device, enter `set pool=` and specify another online pool.

> **Note:**
>
> If you have a single pool, the pool name is not displayed, but it is selected.

```
hostname:configuration storage (pool0)> set pool=pool1
                            pool = pool1
```

3. Enter `show` to see the device information for the pool.

```
hostname:configuration storage (pool1) verify> show
ID STATUS  ALLOCATION  DATA   LOG      CACHE   META     RPM
0      ok     custom      0   0          0/4    0/4    1.86T
1      ok     custom      0   0/2  34G     0      0    15000
2      ok     custom      0   0/2  34G   0/2      0    15000
```

4. Enter `remove`.

```
hostname:configuration storage (pool1)> remove
```

5. Specify the controller or disk shelf, and the number of log or cache devices to remove.

   ID `0` is the controller, and the remaining IDs are the disk shelves. In the following example, `1-log=2` removes two log devices from the first disk shelf:

```
hostname:configuration storage (pool1) remove> set 1-log=2
                            1-log = 2
```

   This example removes one cache device from the controller:

```
hostname:configuration storage (pool1) remove> set 0-cache=1
                            0-cache = 1
```

6. Enter `done`.

```
hostname:configuration storage (pool1) remove> done
```

> **Note:**
>
> If the log devices use a mirrored profile, a message reminds you to select an even number of log devices to remove. If the log devices use a striped profile, you may remove an even or odd number of devices.

Related Topics

Adding a Cache, Meta, or Log Device to an Existing Storage Pool (CLI)

Unconfiguring a Storage Pool (BUI)

Unconfiguring a storage pool removes any active filesystems and LUNs and makes the raw storage available for future storage configuration. This process can be undone by importing the unconfigured storage pool, if the raw storage has not since been used as part of an active storage pool.

Note:

If the storage pool contains a file that is actively retained using the mandatory file retention policy, the pool cannot be unconfigured until all mandatory file retention has expired. For more information, see File Retention Management.

Caution:

Unconfiguring a pool renders data inaccessible, creates the potential for data loss, and fails inbound replications.

Before You Begin

- Do not unconfigure a pool while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** menu, select **System**.
- Do not unconfigure a pool while the peer controller is down or unreachable.
- If an error message reports that the target is in use, wait and try the operation again.

1. From the **Configuration** menu, select **Storage**.
2. From the **Available Pools** list, select an online pool to unconfigure.
3. Click **UNCONFIG**.

Related Topics

- Importing an Existing Storage Pool (BUI)
- Renaming a Storage Pool (BUI)

Unconfiguring a Storage Pool (CLI)

Unconfiguring a storage pool removes any active filesystems and LUNs and makes the raw storage available for future storage configuration. This process can be undone by importing the unconfigured storage pool, if the raw storage has not since been used as part of an active storage pool.

Note:

If the storage pool contains a file that is actively retained using the mandatory file retention policy, the pool cannot be unconfigured until all mandatory file retention has expired. For more information, see File Retention Management.

Caution:

Unconfiguring a pool renders data inaccessible, creates the potential for data loss, and fails inbound replications.

Before You Begin

- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to `maintenance system updates`.
- Do not unconfigure a pool while the peer controller is down or unreachable.
- If an error message reports that the target is in use, wait and try the operation again.

1. Go to `configuration storage`.
2. If you have multiple pools, a default pool is selected and displayed. If this is not the pool you want to unconfigure, enter `set pool=` and specify another online pool.

   > **Note:**
   >
   > If you have a single pool, the pool name is not displayed, but it is selected.

   ```
   hostname:configuration storage (pool0)> set pool=pool1
                                  pool = pool1
   ```

3. Enter `unconfig`.

   ```
   hostname:configuration storage (pool1)> unconfig
   ```

4. Enter `done`.

Related Topics

- Importing an Existing Storage Pool (CLI)
- Renaming a Storage Pool (CLI)

Renaming a Storage Pool (BUI)

To rename a storage pool, you must unconfigure it and then immediately import it with a new name. While storage is unconfigured, data will be inaccessible and there is a potential for data loss. Importing a storage pool can take a considerable amount of time.

Before You Begin

- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check whether an upgrade is in progress, from the **Maintenance** menu, select **System**.
- Do not rename a pool while the peer controller is down or unreachable.

1. From the **Configuration** menu, select **Storage**.
2. From the **Available Pools** list, select the online pool to rename.
3. Click **UNCONFIG**, then **COMMIT**.
4. Click **IMPORT**, then select the storage pool you just unconfigured.
5. Click the storage pool name and change it.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (.) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **- . :**

6. Click **COMMIT**.

Related Topics

- Unconfiguring a Storage Pool (BUI)

- Importing an Existing Storage Pool (BUI)

Renaming a Storage Pool (CLI)

To rename a storage pool, you must unconfigure it and then immediately import it with a new name. While storage is unconfigured, data will be inaccessible and there is a potential for data loss. Importing a storage pool can take a considerable amount of time.

Before You Begin

- Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check whether an upgrade is in progress, navigate to `maintenance system updates`.
- Do not rename a pool while the peer controller is down or unreachable.

1. Go to `configuration storage`.
2. Select the pool that you want to rename.

   If you have multiple pools, a default pool is selected and displayed. If this is not the pool you want to rename, enter `set pool=` and specify another online pool.

   > **Note:**
   >
   > If you have a single pool, the pool name is not displayed, but it is selected.

   ```
   hostname:configuration storage (pool0)> set pool=pool1
                             pool = pool1
   ```

3. Enter `unconfig`.

   ```
   hostname:configuration storage (pool1)> unconfig
   ```

4. Enter `done`.
5. Enter `import`.

   ```
   hostname:configuration storage> import

   Search for storage. Begin the process of searching for existing storage pools.

   Subcommands that are valid in this context:

      help [topic]           => Get context-sensitive help. If [topic] is specified,
                                it must be one of "builtins", "commands", "general",
                                "help" or "script".

      show                   => Show information pertinent to the current context

      abort                  => Abort this task (potentially resulting in a
                                misconfigured system)

      done                   => Finish operating on "discover"

   hostname:configuration storage discover>
   ```

6. Enter `done`.
7. To select the storage pool you just unconfigured, enter `set pool=` and the pool name.

```
hostname:configuration storage select> set pool=pool1
                          pool = pool1
```

8. To rename the storage pool, enter `set name=` and a new name.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (.) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **- . :**

   ```
   hostname:configuration storage (pool1)> set name=NewPool
                             pool = NewPool
   ```

9. Enter `done`.

Related Topics

- Unconfiguring a Storage Pool (CLI)
- Importing an Existing Storage Pool (CLI)

Scrubbing a Storage Pool – Manual (BUI)

Scrubbing a storage pool verifies the content by checking for errors. Scheduled storage pool scrubbing is enabled by default, as described in Scrubbing a Storage Pool – Scheduled (BUI). The recommended minimum period for performing a scrub is quarterly.

- A scrub should be performed at least as often as your oldest backup expires.
- A scrub should also be run before performing a software upgrade.

A scrubbing operation will not proceed if a scrubbing or resilvering is already in progress. If a scrubbing operation is in progress when a resilvering starts, the resilvering operation suspends the current scrubbing and restarts the scrubbing after the resilvering is complete.

1. From the **Configuration** menu, select **Storage**.
2. From the **Available Pools** list, select the online pool to scrub.
3. Optional: Check the status of the last scrub.

   The **Scrub Status** field shows when the last scrub completed in local time (or how long the scrub ran if it did not complete) and the number of errors that were reported by the scrub.

   The `zpool status` command also reports the completion time of the last scrub, how long the scrub ran, and whether errors were found and repaired.

4. Click **SCRUB**.

   The **Scrub Status** field is updated.

5. Optional: To stop the scrub, click **CANCEL**.

   Clicking **SCRUB** again restarts the scrub.

Related Topics

- Storage Pool Concepts
- Scrubbing a Storage Pool – Scheduled (BUI)

Scrubbing a Storage Pool – Manual (CLI)

Scrubbing a storage pool verifies the content by checking for errors. Scheduled storage pool scrubbing is enabled by default, as described in Scrubbing a Storage Pool – Scheduled (CLI). The recommended minimum period for performing a scrub is quarterly.

- A scrub should be performed at least as often as your oldest backup expires.
- A scrub should also be run before performing a software upgrade.

A scrubbing operation will not proceed if a scrubbing or resilvering is already in progress. If a scrubbing operation is in progress when a resilvering starts, the resilvering operation suspends the current scrubbing and restarts the scrubbing after the resilvering is complete.

1. Go to `configuration storage`.
2. Select the online pool to scrub.

   If you have a single pool, the pool name is not displayed, but it is selected.

   If you have multiple pools, one of the pools is selected and displayed. Use the `set pool=` command to select a different pool.

   ```
   hostname:configuration storage (pool0)> set pool=pool1
                              pool = pool1
   ```
3. Optional: Check the status of the last scrub.

   Use the `ls`, `show`, or `get` command to see when the last scrub completed in GMT (or how long it ran if it did not complete), and the number of errors that were reported by the scrub:

   ```
   hostname:configuration storage (pool1)> get scrub
              scrub = scrub completed after 2d20h with 0 errors on Thu May  6 10:35:16
   2022
   ```

   The `zpool status` command also reports the completion time of the last scrub, how long the scrub ran, and whether errors were found and repaired.
4. Enter `scrub start`.

   ```
   hostname:configuration storage (pool1)> scrub start
   ```
5. Optional: Stop the scrub before it has completed by entering `scrub stop`.

   ```
   hostname:configuration storage (pool1)> scrub stop
   ```

   Entering `scrub start` again restarts the scrub.

Related Topics

- Storage Pool Concepts
- Scrubbing a Storage Pool – Scheduled (CLI)

Scrubbing a Storage Pool – Scheduled (BUI)

Scrubbing a storage pool verifies the content by checking for errors. By default, scheduled storage pool scrubbing is enabled and set to every 30 days. Use the **Scrub Schedule** control to specify a different scrub interval or to disable scrub scheduling.

You cannot change the scrub priority explicitly. Scrub priority is automatically adjusted based on the specified scrub interval, the progress of the scrub, and the system load. Scrub priority automatically increases on an idle system.

A scrubbing operation will not proceed if a scrubbing or resilvering is already in progress. If a scrubbing operation is in progress when a resilvering starts, the resilvering operation suspends the current scrubbing and restarts the scrubbing after the resilvering is complete.

1. From the **Configuration** menu, select **Storage**.
2. From the **Available Pools** list, select the online pool to configure for scheduled scrubbing.

3. Optional: Check the status of the last scrub.

   The **Scrub Status** field shows when the last scrub completed in local time (or how long the scrub ran if it did not complete), and the number of errors that were reported by the scrub.

   The `zpool status` command also reports the completion time of the last scrub, how long the scrub ran, and whether errors were found and repaired.

4. Select **Scrub Schedule**.

   Choices on the **Scrub Schedule** drop-down menu are:

   - **Off**
   - **15 days**
   - **30 days**
   - **45 days**
   - **60 days**
   - **75 days**
   - **90 days**

   By default, scrub is scheduled for every **30 days**. You can select a different scrub interval or, to disable scheduled scrubbing, select **Off**.

5. Click the **SCHEDULE** button.

Related Topics

- Storage Pool Concepts
- Scrubbing a Storage Pool – Manual (BUI)

Scrubbing a Storage Pool – Scheduled (CLI)

Scrubbing a storage pool verifies the content by checking for errors. By default, scheduled storage pool scrubbing is enabled and set to every 30 days. Use the `scrub_schedule` property to specify a different scrub interval or to disable scrub scheduling.

You cannot change the scrub priority explicitly. Scrub priority is automatically adjusted based on the specified scrub interval, the progress of the scrub, and the system load. Scrub priority automatically increases on an idle system.

A scrubbing operation will not proceed if a scrubbing or resilvering is already in progress. If a scrubbing operation is in progress when a resilvering starts, the resilvering operation suspends the current scrubbing and restarts the scrubbing after the resilvering is complete.

1. Go to `configuration storage`.

2. Select the online pool to scrub as follows:

   - If you have a single pool, the pool name is not displayed, but it is selected.
   - If you have multiple pools, one of the pools is selected and displayed. Enter the `set pool=` command to specify a different pool.

   ```
   hostname:configuration storage (pool0)> set pool=pool1
                                 pool = pool1
   ```

3. Optional: Check the status of the last scrub.

   Use the `ls`, `show`, or `get` command to see when the last scrub completed in GMT (or how long it ran if it did not complete), and the number of errors that were reported by the scrub:

```
hostname:configuration storage (pool1)> get scrub scrub_schedule
           scrub = scrub completed after 2d20h with 0 errors on Thu May  6 10:35:16
2022
  scrub_schedule = 30 days
```

The `zpool status` command also reports the completion time of the last scrub, how long the scrub ran, and whether errors were found and repaired.

4. Set a value for the `scrub_schedule` property.

   Set the value of the `scrub_schedule` property to `off` to disable scrub scheduling, for example if you prefer to perform a manual scrub.

   To enable scheduled storage pool scrubbing, set the value of the `scrub_schedule` property to the number of days between scheduled scrubs. You can set `scrub_schedule` to 15, 30, 45, 60, 75, or 90 days.

```
hostname:configuration storage (p0)> set scrub_schedule=60
                          scrub_schedule = 60
hostname:configuration storage (p0)> get scrub_schedule
                          scrub_schedule = 60 days
```

Related Topics

- Storage Pool Concepts
- Scrubbing a Storage Pool – Manual (CLI)

Viewing Pool and Device Status (BUI)

You can check the status of pool and component devices. If there is a problem with a pool, details about device status will also be listed.

1. From the **Configuration** menu, select **Storage**.
2. Click on a pool to select it and see more details.

   Refer to the following table for a description of the pool status.

Pool Status	Description
Online	A pool that has all devices operating normally.
Degraded	A pool with one or more failed devices, but the data is still available due to a redundant configuration.
Faulted	One or more component devices are offline and there are insufficient replicas to continue functioning.
Offline	A pool was explicitly taken offline.
Unavailable	A pool with corrupted metadata, or one or more unavailable devices and insufficient replicas to continue functioning.
Exported	A pool is active on the cluster peer and is ready for a cluster failback to occur.
Retention	A pool has at least one actively retained file with the mandatory file retention policy. The retention status is also listed with the timestamp of the last or latest file retained within any of the pool's filesystems with mandatory retention, and "(expired)" is displayed if that time has passed. The number of filesystems with mandatory file retention and their files is also displayed.

3. View pool encryption properties.

   - If the appliance is at software release OS8.8.0 or later and the "Enable Pool Encryption" deferred update is applied, an **ENCRYPTED** column is displayed. For more information, see Enable Pool Encryption Deferred Update in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

- If a pool is not encrypted, then the **ENCRYPTED** column is blank.
- If a pool is encrypted, then a lock icon 🔒 is displayed in the **ENCRYPTED** column, and hovering the mouse pointer over the lock icon displays the encryption type.
- If the selected pool is encrypted, then information about the encryption type, keystore, and key name is shown. See Encryption Properties.
  - If the value of the **Key Status** field is **unavailable**, then the key has been deleted. In addition, an indicator ● is displayed to the right of the lock icon 🔒 and shows that the key for this pool is deleted.
  - If the **Key Last Update** field is blank, then the key value has not been changed since the pool was created.

4. View the selected pool's device statuses under the **Device Status** section.

   Refer to the following table for a description of the device status.

Device Status	Description
Online	The device is online and functioning. You may not see this status, and instead see the message "No device faults have been detected in the storage pool."
Degraded	The device is not in an optimal state. Either it is expected to fail soon, or a spare has not finished resilvering yet.
Faulted	The device is faulty; more information can be found in the maintenance logs.
Offline	The device was explicitly taken offline; no reads or writes will occur to this device until it has been onlined.
Removed	The device has been physically removed.
Hot Spare	This spare device is actively being used as a data disk in the pool as a replacement for a device that failed.
Unavailable	The device could not be opened or the pool could not detect this device.

5. To see more detailed pool and device error information, do the following:
   - To view active errors, from the **Maintenance** menu, select **Problems**.
   - To view the history for all problems, from the **Maintenance** menu, select **Logs**.

Viewing Pool and Device Status (CLI)

Use the following procedure to view pool and device status.

1. Go to `configuration storage`.

   The default pool is selected and displayed.

2. Optional: If you have multiple pools, select a different pool to view.

   If the default pool is not the pool that you want to view, enter `set pool=` and specify another online pool.

```
hostname:configuration storage (pool0)> set pool=pool1
                         pool = pool1
```

3. Enter `show`.
   - See the table in step 2 of Viewing Pool and Device Status (BUI) for a description of the pool status value.

- See the table in step 4 of Viewing Pool and Device Status (BUI) for a description of the device status value.

4. View pool encryption properties.

    - If the appliance is at software release OS8.8.0 or later and the Enable Pool Encryption deferred update is applied (see Enable Pool Encryption Deferred Update in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*), an `ENCRYPTED` column is displayed.
    - If a pool is not encrypted, then the value of the `ENCRYPTED` column is a hyphen (`-`) and the value of the `encryption` property is `off`.
    - If a pool is encrypted, then the `ENCRYPTED` column displays the encryption type. See Encryption Properties.
    - If the selected pool is encrypted, then the `encryption` property displays the encryption type. In addition, the following four properties are displayed:
        - `keystatus` - Either `available` or `unavailable`. If the value of the `keystatus` property is `unavailable`, then the key has been deleted.
        - `keychangedate` - The date the key was last changed. If the value of the `keychangedate` property is blank, then the key value has not been changed since the pool was created.
        - `keyname` - The name of the key.
        - `keystore` - The name of the keystore.

```
hostname:configuration storage (pool1)> ls
Pools:
  POOL            OWNER     DATA PROFILE LOG PROFILE STATUS   ERRS ENCRYPTED
->pool1           hostname  mirror       -           online   0    aes-128-ccm
  pool2           hostname  mirror       -           online   0    -
  pool3           hostname  stripe       -           online   0    -

Properties:
                          pool = pool1
                        status = online
                        errors = 0
                       profile = mirror
                   log_profile = -
                 cache_profile = -
                  meta_profile = -
            retainedfileystems = 1
               retentionexpiry = 2022-05-16-20:20:35 (expired)
                         scrub = none requested
   async_destroy_reclaim_space = 0
                    encryption = aes-128-ccm
                     keystatus = available
                 keychangedate =
                       keyname = MyKey
                      keystore = LOCAL
```

If a pool is not encrypted, the four `key` values do not display but are still available by using the `get` or `script` commands.

5. To see more detailed pool and device error information, go to `maintenance problems` for active errors, or `maintenance logs` for a history of all problems.

Storage Pool Concepts

Storage is configured in pools that are characterized by their underlying data redundancy, and provide space that is shared across all filesystems and LUNs. More information about how storage pools relate to individual filesystems or LUNs can be found in About Storage Pools, Projects, and Shares.

To understand storage pool concepts, use these topics:

- Storage Pool Configuration
- Multiple Pools
- Number of Devices per Pool
- Drive Characteristics and Performance
- Storage Pool Capacity
- All-Flash Storage Configuration
- Storage Pool Reclaimed Space
- Destroy Prevention and Approval

Related Topics

- Configuring an All-Flash Storage Pool - BUI, CLI
- Disk Shelf Configurations in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*
- Monitoring SSD Endurance in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*

Storage Pool Configuration

Pools can be created by configuring a new pool or importing an existing pool. Importing an existing pool is only used to import pools that were previously configured on Oracle ZFS Storage Appliance. Importing an existing pool is useful in case of accidental reconfiguration, such as when moving pools between controllers, or due to catastrophic controller failure.

Multiple Pools

Each controller can have any number of pools, and each pool can be assigned ownership independently in a cluster. With the ability to control access to log and cache devices on a per-share basis, the recommended mode of operation is a single pool. While arbitrary number of pools are supported, creating multiple pools with the same redundancy characteristics owned by the same cluster head is not advised. Doing so will result in poor performance, suboptimal allocation of resources, artificial partitioning of storage, and additional administrative complexity. Configuring multiple pools on the same host is only recommended when drastically different redundancy or performance characteristics are desired, for example a mirrored pool for databases and a RAID-Z pool for streaming workloads.

Number of Devices per Pool

Drives within all of the chassis can be allocated individually; however, care should be taken when allocating disks from disk shelves to ensure optimal pool configurations. In general, fewer pools with more disks per pool are preferred because they simplify management and provide a higher percentage of overall usable capacity.

While the system can allocate storage in any increment desired, it is recommended that each allocation include a minimum of 8 disks across all disk shelves and ideally many more.

Drive Characteristics and Performance

Follow these restrictions when configuring storage pools:

- All data disks contained within a head node or disk shelf must have the same rotational speed (media rotation rate). The appliance software will detect misconfigurations and generate a fault for the condition.
- Due to unpredictable performance issues, avoid mixing different disk rotational speeds within the same pool.
- For optimal performance, do not combine disk shelves with different disk rotational speeds on the same SAS fabric (HBA connection). Such a mixture operates correctly, but likely results in slower performance of the faster devices.
- When creating a new pool, avoid mixing different data disk capacities because all disks are then limited to the smallest capacity disk in the pool. When adding a higher capacity disk to an existing pool, the larger disk's capacity is maintained. However, the system preferentially writes to new disks until they begin to reach the same capacity utilization as the old disks. To maintain performance, add as many new higher capacity disks as the total number of disks in the original pool.
- A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 or later.

Storage Pool Capacity

When allocating raw storage to pools, keep in mind that filling pools completely will result in significantly reduced performance, especially when writing to shares or LUNs. These effects become more noticeable as the pool reaches full capacity.

All-Flash Storage Configuration

Oracle Storage Drive Enclosure DE3-24P can be configured as all-flash storage with fully populated flash-based SSD data devices and optional log devices. All-flash storage provides low-latency I/O that increases workload performance.

An all-flash storage pool contains data SSDs and optional log devices. Read flash cache and meta devices cannot be part of an all-flash pool. The remaining lifetime of SSDs can be monitored using threshold alerts.

Storage Pool Reclaimed Space

When deleting a project, filesystem, or LUN, you can view the amount of space to be reclaimed in the storage pool if deferred update Asynchronous Dataset Deletion (OS8.7.0 or later) has been accepted. In the BUI, field **Asynchronous Dataset Destroy** is displayed during these deletion operations. Similarly in the CLI, property `async_destroy_reclaim_space` reflects the amount of space to be reclaimed and shows `0` (zero) when the operation has completed. The individual procedures to delete a project, filesystem, or LUN contain a step for monitoring the reclaimed space in a storage pool.

Destroy Prevention and Approval

When the **prevent share destruction** (`nodestroy`) property is set to on/true at the storage pool level, shares and projects within the pool cannot be deleted. Replication packages within the pool, and the pool itself, are also protected from destruction because of inheritance rules. Protection is also extended to destroying a share through dependent clones. However, it does not affect shares destroyed through replication updates. If a share is destroyed on an Oracle ZFS Storage Appliance system that is the source for replication, the corresponding share on the target will be destroyed, even if this property is set to on/true. The property's default setting is off/false.

The **prevent destruction** (`nodestroy`) property is also available at the project level, and all shares within the project inherit this property setting. Shares and projects can be deleted by setting this property to off/false at the share or project level. Similarly, when property **prevent share destruction** (`nodestroy`) is set to on/true at the pool level, all shares, projects, and replication packages in the pool, and the pool itself, are protected. Furthermore, if additional property **destroy requires approver** (`approve_destroy`) is set to on/true, then the delete process requires an extra step. This property is only available at the pool level.

When property **prevent share destruction** (`nodestroy`) is set to on/true at the pool level, and not set at the share or project level, the pool-level property's setting is the default setting.

To delete a share, project, or replication package (CLI `destroy` command) when it is protected at the pool level, two different administrators must perform the delete action: One administrator, the approver, sets the **prevent destruction** (`nodestroy`) property to off/false at the share or project level. A different administrator deletes the share, project, or replication package in the normal manner. The approving administrator's username is recorded in read-only property **destroy approved by** (`destroy_approved_by`). In the BUI, the approver's username is displayed in the project's sidebar, along with other static properties, as well as on the pool's properties page. If all projects within the storage pool have been approved for deletion, then the pool itself can be unconfigured.

When the **prevent destruction** (`nodestroy`) property is set again to on/true at the project level for the same project, it clears property **destroy approved by** (`destroy_approved_by`).

> **Note:**
>
> These destruction-prevention methods do not apply to snapshots. For snapshots, see the following tasks listed in Snapshots and Clones:
>
> - Viewing Snapshots and Schedules - BUI, CLI
> - Renaming a Snapshot - BUI, CLI
> - Editing a Snapshot Retention Policy - BUI, CLI

Related Topics:

- Project Properties
- Filesystem Properties
- LUN Properties

Data Profiles for Storage Pools

After storage devices are physically verified and resources are allocated for a storage pool, the next step is to choose a storage profile that reflects your reliability, availability, serviceability

(RAS), and performance goals. The set of possible profiles presented depends on your available storage. The following table lists all possible profiles and their descriptions.

Table 2-11 Data Profiles

Data Profile	Description
Dual Parity Options	
Triple mirrored	Data is triply mirrored, yielding a very highly reliable and high-performing system (for example, storage for a critical database). This configuration is intended for situations in which maximum performance and availability are required. Compared with a two-way mirror, a three-way mirror adds additional IOPS per stored block and higher level protection against failures. Note: A controller without expansion storage should not be configured with triple mirroring.
Double parity RAID	RAID in which each stripe contains two parity disks. As with triple mirroring, this yields high availability, as data remains available with the failure of any two disks. Double parity RAID is a higher capacity option than the mirroring options and is intended either for high-throughput sequential-access workloads (such as backup) or for storing large amounts of data with low random-read component.
Single Parity Options	
Mirrored	Data is mirrored, reducing capacity by half, but yielding a highly reliable and high-performing system. Recommended when space is considered ample, but performance is at a premium (for example, database storage).
Single parity RAID, narrow stripes	RAID in which each stripe is kept to three data disks and a single parity disk. For situations in which single parity protection is acceptable, single parity RAID offers a much higher capacity option than simple mirroring. This higher capacity needs to be balanced against a lower random read capability than mirrored options. Single parity RAID can be considered for non-critical applications with a moderate random read component. For pure streaming workloads, give preference to the double parity RAID option which has higher capacity and more throughput.
Other	
Striped	Data is striped across disks, with no redundancy. While this maximizes both performance and capacity, a single disk failure will result in data loss. This configuration is not recommended. For pure streaming workloads, consider using double parity RAID. Due to non-redundancy, disks configured in a striped profile will not receive firmware updates, unless the configured storage pools are in an exported state.
Triple parity RAID, wide stripes	RAID in which each stripe has three disks for parity. This is the highest capacity option apart from Striped Data. Resilvering data after one or more drive failures can take significantly longer due to the wide stripes and low random I/O performance. As with other RAID configurations, the presence of cache can mitigate the effects on read performance. This configuration is not generally recommended.

> **Note:**
>
> Earlier software versions supported double parity with wide stripes. This has been supplanted by triple parity with wide stripes, as it adds significantly better reliability. Pools configured as double parity with wide stripes under a previous software version continue to be supported, but newly configured or reconfigured pools cannot select that option.

NSPF Option

For expandable systems, some profiles may be available with an "NSPF" option. This stands for "no single point of failure" and indicates that data is arranged in mirrors or RAID stripes such that a pathological disk shelf failure will not result in data loss. Note that systems are already configured with redundancy across nearly all components. Each disk shelf has redundant paths, redundant controllers, and redundant power supplies and fans. The only failure that NSPF protects against is disk backplane failure (a mostly passive component), or gross administrative misconduct (detaching both paths to one disk shelf). There are unique cases where NSPF could result in lower capacity, such as when disk trays have a different number of data disks.

Log Devices

Log devices can be configured using only striped or mirrored profiles. Log devices are only used in the event of node failure. For data to be lost with unmirrored logs, it is necessary for both the device to fail and the node to reboot immediately after. This a highly unlikely event; however, mirroring log devices can make this effectively impossible, requiring two simultaneous device failures and node failure within a very small time window.

> **Note:**
>
> When different sized log devices are in different chassis, only striped log profiles can be created.

Cache Devices

In a cluster configuration, cache devices installed in controller slots are available only to the controller which has the storage pool imported. In a cluster, it is possible to configure cache devices on both controllers to be part of the same pool. To do this, take over the pool on the passive node, then add storage, and select the cache devices. This has the effect of having half the global cache devices configured at any one time. While the data on the cache devices will be lost on failover, the new cache devices can be used on the new controller.

Cache devices installed in disk shelf slots, when added to a pool, are automatically imported during a cluster failback or takeover. No additional pool configuration is required.

Meta Devices

A meta device is a cache device used to store deduplicated metadata and other metadata for projects and shares. Meta devices can be allocated to a storage pool, but not an all-flash storage pool, during and after storage pool creation. However, they cannot be reconfigured as normal cache devices for a pool, nor can they be removed from a pool. A meta device must be a 3.2 TB (minimum) SSD to support the enhanced data deduplication feature available in software version OS8.7.0 (2013.1.7.0) or later.

Before using meta devices and the deduplication feature for new and existing storage pools, accept the deferred software update for Data Deduplication v2, introduced with software version OS8.7.0 (2013.1.7.0). If replicating to other systems, both the replication source and targets must have this deferred update. For more information, see Data Deduplication, and Data Deduplication v2 Deferred Update in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

Hot Spares

Hot spares are allocated as a percentage of total pool size and are independent of the profile chosen (with the exception of striped, which does not support hot spares). Because hot spares are allocated for each storage configuration step, it is much more efficient to configure storage as a whole than it is to add storage in small increments.

Related Topics:

- Creating a Storage Pool - BUI, CLI
- Adding a Cache, Meta, or Log Device to an Existing Storage Pool - BUI, CLI

Configuring Cloud Backup

Cloud backups provide the ability to back up full and incremental share snapshots from a local Oracle ZFS Storage Appliance to cloud targets in an Oracle Cloud Infrastructure account. Share snapshots are created as usual, and individual snapshots are selected for cloud backup. More than one appliance can back up to the same cloud target, and cloud backups can be restored on any Oracle ZFS Storage Appliance with access to the cloud target.

By using the "tar" archiving format for cloud backups, as opposed to the default "zfs" format, you can create full and incremental filesystem (not LUN) snapshot backups. These backups can be restored on any system, regardless of its operating system, that can access the cloud target. However, a tar cloud backup does not preserve the original filesystem properties.

After you have established an Oracle Cloud Infrastructure account and created buckets in Object Storage, you prepare Oracle ZFS Storage Appliance systems to use cloud backup, define cloud targets, and create cloud backups that can later be restored.

In the BUI, there are five tabs for the cloud data service: **Properties**, **Targets**, **Backups**, **Jobs**, and **Logs** (service log). In the CLI, the cloud data service has three child nodes: `backups`, `targets`, and `jobs`. Objects in each of these CLI child nodes are assigned ordinal names, such as `backup-001`. In either the BUI or CLI, creating a cloud backup creates a cloud job with an operation of `backup`. The other operations for cloud jobs are `restore` and `delete`. Like other services, user operations are logged to the audit log, and results of operations are logged to the alert log. Cloud service state changes are logged to the service log.

To manage cloud storage, you can view the cloud jobs, cancel and restart jobs, modify or delete cloud targets, view and filter cloud backups, and restore and delete cloud backups. To analyze backup and restore operations, see Data Movement: Cloud Bytes and Data Movement: Cloud Requests in *Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x*. To set alerts, see Configuring Alerts.

Before using cloud backup on clustered controllers, it is important to understand clustering behavior in conjunction with this feature. The cloud data service is cluster aware, and each controller uses its own available network interfaces to communicate with configured cloud targets to perform backup and restore operations, and to refresh its list of available cloud backups. Therefore, a controller might not immediately have a refreshed list of cloud backups from its peer controller. Also, while both controllers share common service configurations, cloud targets, and cloud backups, each controller has its own Jobs list. After a restart or failover event, cloud backup or restore operations running on the impacted controller are automatically resubmitted on the peer controller, and the operations restart from the beginning.

This section contains the following topics and tasks:

- Preparing for Cloud Backup
  - Setting Up Your Oracle Cloud Infrastructure Account
  - Adding Cloud Authorizations to a User Role
  - Enabling the Cloud Data Service - BUI, CLI
- Configuring Cloud Backup and Restore Operations
  - Defining a Cloud Target - BUI, CLI
  - Creating a Cloud Backup - BUI, CLI
  - Viewing and Filtering Cloud Backups - BUI, CLI
  - Restoring a Cloud Backup - BUI, CLI
- Managing the Cloud Service
  - Deleting a Cloud Backup - BUI, CLI
  - Viewing Cloud Jobs - BUI, CLI
  - Canceling and Restarting a Cloud Job - BUI, CLI
  - Modifying a Cloud Target - BUI, CLI
  - Deleting a Cloud Target - BUI, CLI
- Cloud Properties

Preparing for Cloud Backup

To prepare for cloud backup, perform the following tasks:

- Setting Up Your Oracle Cloud Infrastructure Account
- Adding Cloud Authorizations to a User Role
- Enabling the Cloud Data Service - BUI, CLI

Setting Up Your Oracle Cloud Infrastructure Account

Oracle Cloud Infrastructure offers two distinct storage tiers that are robust and scalable: Standard storage is for frequently accessed storage, and archive storage is for less frequently accessed storage, requires a longer time to retrieve the data, and is more cost effective. Thus, the two tiers allow for maximum performance versus minimum cost. Oracle Cloud Infrastructure Object Storage addresses storage, disaster recovery, data migration, and master image replication needs, while also providing data encryption during transport and at rest.

Create an Oracle Cloud Infrastructure account and configure it, including creating buckets for your data. A bucket's tier can either be "standard" or "archive", and you cannot change a bucket's tier after creating the bucket. The cloud backup stores the data, the share snapshot; the metadata stores information about the data. Optionally, you can store metadata and data in different buckets. If you use an archive bucket, the metadata must be saved to a standard bucket.

Each cloud target uses one regional API endpoint (location), and each bucket can be used by multiple appliances. If you need to create multiple cloud targets with the same endpoint, create enough distinct buckets to contain them. From the Oracle Cloud Infrastructure Console, be sure to note:

- Regional API endpoint URLs, such as `https://objectstorage.us-phoenix-1.oraclecloud.com`
- Tenancy name and Oracle Cloud Identifier (OCID)

  For the OCID, go to your compartment. In the **Tenancy Information** tab, click **Show** or **Copy** for the OCID.
- User name and OCID

  For the OCID, select your user name from the user icon in the upper right-hand corner. In the **User Information** tab, click **Show** or **Copy** for the OCID.
- Private key: Create an RSA public/private key pair. The key must be in the Privacy Enhanced Mail (PEM) format. For more information, see "Required Keys and OCIDs" in Oracle Cloud Infrastructure API documentation available at https://docs.oracle.com/en-us/iaas/Content/API/Concepts/apisigningkey.htm .
- Bucket names

For information about the Oracle Cloud Infrastructure Object Storage Service API, see “Object Storage Service API” in Oracle Cloud Infrastructure API documentation available at https://docs.oracle.com/en-us/iaas/api/#/en/objectstorage/20160918/ . Here you will find such information as regional API endpoints (locations), how to create and manage buckets, and optional object lifecycle management.

> ⚠ **Caution:**
>
> If you alter cloud backup objects in buckets in Oracle Cloud Infrastructure Object Storage, you will create inconsistencies with content stored locally. Therefore, it is recommended to use dedicated buckets for cloud backup.

Adding Cloud Authorizations to a User Role

For the appropriate user role(s) and if all services are not globally authorized, add authorizations for the **Cloud** filter for the **Services** scope. Like other services, there are three available authorizations: **administer** (enable and disable the service), **configure** (change service settings), and **restart** (restart the service).

Also, add scope **Cloud Targets**, select all cloud targets or a specific cloud target, and select authorizations:

- **backup** - Back up snapshot data to cloud
- **delete** - Delete cloud backups
- **restore** - Restore cloud backups to local shares

For scope **Projects and Shares**, globally set filters or specify filters, and add authorizations:

- **backup** - Read share data to back up to cloud target if cloud **backup** is authorized for the cloud target
- **restore** - Write cloud backup data if cloud **restore** is authorized for the cloud target

Use the appropriate user's role with the corresponding required authorizations for configuring and managing the cloud data service and cloud backups. For more information about user configuration, see Configuring Users.

Enabling the Cloud Data Service (BUI)

Use the following procedure to enable the cloud data service.

1. From the **Configuration** menu, select **Services**.
2. Under **Data Services**, click the power icon ⏻ for **Cloud** to enable the service.

   The status icon turns green.

Enabling the Cloud Data Service (CLI)

Use the following procedure to enable the cloud data service.

1. Go to `configuration services cloud`.

   ```
   hostname:> configuration services cloud
   ```

2. Enable the cloud service.

   Enter `get` to determine if the cloud service is enabled, and has a status of `online`. If the service is disabled, enter `enable`.

   ```
   hostname:configuration services cloud> get
   Properties:
               <status> = disabled
   .
   .
   .
   hostname:configuration services cloud> enable
   ```

3. To confirm that the cloud service status is online, enter `get`.

   ```
   hostname:configuration services cloud> get
   Properties:
               <status> = online
   ```

Configuring Cloud Backup and Restore Operations

To configure cloud backup and restore operations, use these tasks:

- Defining a Cloud Target - BUI, CLI
- Creating a Cloud Backup - BUI, CLI
- Viewing and Filtering Cloud Backups - BUI, CLI
- Restoring a Cloud Backup - BUI, CLI

Defining a Cloud Target (BUI)

After enabling the cloud data service, you can create a cloud target, which is the combination of location, bucket, and data bucket. Each cloud target uses HTTPS certificates assigned to the cloud data service from the system's trust chain and any other uploaded trusted certificates assigned to the service. For information about assigning a certificate to the cloud data service, see Assigning a Certificate to a Service (BUI).

1. From the **Configuration** menu, select **Services**.

2. Under **Data Services**, click **Cloud**.
3. Optional: To set the TLS version and ciphers, go to tab **Properties**, click the edit icon ✎ for **SSL/TLS Versions and Ciphers**, and complete the dialog box.
4. Go to tab **Targets**, and click the add icon ⊕ next to **Targets**.
5. In the **Add Cloud Target** dialog box, enter the following configuration information, and click **ADD**.
   - **Name** - Your name for this cloud target, which must be unique for your system.
   - **Location** - The URL for a regional API endpoint (location), such as `https://objectstorage.us-phoenix-1.oraclecloud.com`
   - **Use data bucket** - Select this check box if your data and metadata will be in different buckets. Data will be saved to the `Data Bucket`, and metadata will be saved to the `Bucket`. If you use an archive bucket, the metadata must be saved to a standard bucket.
   - **Bucket** - An established bucket name. Data can be saved in either a standard or archive bucket. The combination of `Location`, `Bucket`, and `Data Bucket` must be unique for your system. If you selected `Use data bucket`, this field is for the metadata bucket; if you did not select `Use data bucket`, this field is for both the metadata and data. To specify a namespace, include it in the bucket name, such as `namespace/bucket`; for example, `export/test2/bucket` where `export/test2` is the namespace. This can be used to specify a different share name for the target appliance, with the namespace used as the mountpoint.
   - **Data Bucket** - If you selected `Use data bucket`, this field is for the data bucket; if you did not select `Use data bucket`, this field is not available. The combination of `Location`, `Bucket`, and `Data Bucket` must be unique for your system.
   - **User** - Your user name OCID from the Oracle Cloud Infrastructure account.
   - **Tenancy** - Your tenancy name OCID from the Oracle Cloud Infrastructure account.
   - **Private Key** - Your private key for the account, which must match the public key uploaded to the Oracle Cloud Infrastructure account. Enter the key in the Privacy Enhanced Mail (PEM) format.
   - **Use web proxy** - Select this check box to use a proxy for system communications with the web, and complete the fields for host port name and number. The user name and password fields are optional.
   - **Write limit bandwidth** (optional) - Select this check box to limit the traffic write bandwidth when uploading a cloud backup to the cloud target. Enter a value and select a unit of measurement. For example, `5 M/s` limits writes to the cloud target to 5 megabytes per second.
   - **Read limit bandwidth** (optional) - Select this check box to limit the traffic read bandwidth when restoring a cloud backup from the cloud target. Enter a value and select a unit of measurement. For example, `4 M/s` limits reads from the cloud target to 4 megabytes per second.

Related Topics

- Modifying a Cloud Target (BUI)
- Deleting a Cloud Target (BUI)
- Cloud Properties

Defining a Cloud Target (CLI)

After enabling the cloud data service, you can create a cloud target, which is the combination of location, bucket, and data bucket. Each cloud target uses HTTPS certificates assigned to the cloud data service from the system's trust chain and any other uploaded trusted certificates assigned to the service. For information about assigning a certificate to the cloud data service, see Assigning a Certificate to a Service (CLI).

1. Go to `configuration services cloud`.

   ```
   hostname:> configuration services cloud
   ```

2. Optional: To set the TLS versions and ciphers, enter `ls` to list the properties, and then set the `tls_version` and `ciphers` properties.
3. Enter `targets`.
4. List the properties to configure by entering `get`.

   ```
   hostname:configuration services cloud targets> get
   Properties:
                         name = (unset)
                     location = (unset)
                         user = (unset)
                databucket_on = false
                       bucket = (unset)
                  data_bucket = (unset)
                      tenancy = (unset)
                          key = (unset)
                     proxy_on = false
                   proxy_host = (unset)
                   proxy_user = (unset)
               proxy_password = (unset)
                   writelimit = (unset)
                    readlimit = (unset)
   ```

5. To modify the properties, use the `set` command, such as `set name=oci-phoenix`. The configuration properties are as follows:
   - `name` - Your name for this cloud target, which must be unique for your system.
   - `location` - The URL for a regional API endpoint (location), such as `https://objectstorage.us-phoenix-1.oraclecloud.com`
   - `user` - Your user name OCID from the Oracle Cloud Infrastructure account.
   - `databucket_on` - Set this to `true` if your data and metadata will be in different buckets. Data will be saved to the `data_bucket`, and metadata will be saved to the `bucket`. If you use an archive bucket, the metadata must be saved to a standard bucket.
   - `bucket` - An established bucket name. Data can be saved in either a standard or archive bucket. The combination of `location`, `bucket`, and `data_bucket` must be unique for your system. If you set `databucket_on` to `true`, this property is for the metadata bucket; if you set `databucket_on` to `false`, this property is for both the metadata and data. To specify a namespace, include it in the bucket name, such as `namespace/bucket`; for example, `export/test2/bucket` where `export/test2` is the namespace. This can be used to specify a different share name for the target appliance, with the namespace used as the mountpoint.

- `data_bucket` - If you set `databucket_on` to `true`, this property is for the data bucket; if you set `databucket_on` to `false`, this property cannot be set. The combination of `location`, `bucket`, and `data_bucket` must be unique for your system.
- `tenancy` - Your tenancy name OCID from the Oracle Cloud Infrastructure account.
- `key` - Your private key for the account, which must match the public key uploaded to the Oracle Cloud Infrastructure account. Enter the key in the Privacy Enhanced Mail (PEM) format.

  Use the `setkey` command to enter the key in an interactive mode, thus allowing multiple lines. When finished entering all lines, enter a period (".").

  ```
  hostname:configuration services cloud targets> setkey
  ("." to end)> - - - - -BEGIN RSA PRIVATE KEY- - - - -
  ("." to end)>
  MD0CAQACCQDJCnSbjr7nUQIDAQABAggQMmdyfF9wgQIFAOQayRsCBQDhoGQDAgRr
  ("." to end)> NWanAGQeHthbAgQeJsyk
  ("." to end)> - - - - -END RSA PRIVATE KEY- - - - -
  ("." to end)> .
  ```

- `proxy_on` - Set this to `true` to use a proxy for system communications with the web, and complete the property for `proxy_host` (host port name and number). The `proxy_user` (user name) and `proxy_password` (user password) properties are optional. Example for `proxy_host`: `www-proxy.us.example.com:80`.
- `writelimit` (optional) - Set to a value and unit of measurement to limit the traffic write bandwidth when uploading a cloud backup to the cloud target. For example, `5M/s` limits writes to the cloud target to 5 megabytes per second.
- `readlimit` (optional) - Set to a value and unit of measurement to limit the traffic read bandwidth when restoring a cloud backup from the cloud target. For example, `4M/s` limits reads from the cloud target to 4 megabytes per second.

6. When done, enter `commit`.

   ```
   hostname:configuration services cloud (uncommitted)> commit
   ```

Related Topics

- Modifying a Cloud Target (CLI)
- Deleting a Cloud Target (CLI)
- Cloud Properties

Creating a Cloud Backup (BUI)

After defining cloud targets, you can create a cloud backup of a share's (filesystem's or LUN's) full or incremental snapshot to a cloud target(s). Depending on how an incremental cloud backup is created, the parent snapshot must be available on both the local system and backed up to the same cloud target, or the parent snapshot does not also have to be available on the same cloud target. An incremental snapshot backup only backs up the difference between the previous snapshot and the current snapshot. Also, snapshot retention holds are preserved when moving the snapshot to another system via cloud snapshot backup.

Creating a cloud backup creates a cloud job with an operation of `backup`. Snapshots on different systems can have the same name, such as `snap2`, because snapshots are assigned a unique identification for each system.

There are two formats for cloud backups: "zfs" and "tar". The same snapshot on the same system can be used in two cloud backups: one with the zfs format, and one with the tar format. The formats support the features shown in the following table.

Feature	ZFS Format	Tar Format
Restore on only Oracle ZFS Storage Appliance or an Oracle Solaris server	✓	
Restore on any system regardless of operating system		✓
Supports both filesystem and LUN snapshots	✓	
Supports filesystem snapshots only		✓
Preserves filesystem properties	✓	
Preserves LUN properties	✓	
Supports full and incremental backups	✓	✓
Supported within the same Oracle Cloud Infrastructure bucket	✓	✓
High-efficiency compression: • If the underlying share is compressed, less data is transferred and, therefore, the backup is faster. • Incremental backups are performed at the block level.	✓	
Mid-efficiency compression: • Files are read and compressed during the backup operation. • Incremental backups are performed at the file level. Especially not efficient if a large file is modified because the full file will be part of the incremental backup.		✓

> **Note:**
>
> After a local snapshot is in a cloud backup in your Oracle Cloud Infrastructure account, you can delete it locally, per your policies. However, retain local snapshots that could be parent snapshots for future incremental snapshots.

Filesystem snapshots that contain files created with the file retention feature (not the snapshot retention feature) are subject to the rules governing file retention. For information, see Planning Guidelines for File Retention.

1. Ensure the cloud target(s) is online.
   a. From the **Configuration** menu, select **Services**, then **Cloud**.
   b. Click the **Targets** tab.
   c. For each cloud target to be used, ensure its status icon is green. If the status icon is not green, but the cloud data service icon is green, verify that the cloud target properties are valid. See Defining a Cloud Target (BUI).
2. From the **Shares** menu, select **Shares**, then **Filesystems** or **LUNs**, as appropriate.
3. Hover over the share to be backed up, and click its edit icon ✎ .
4. Click the **Snapshots** tab.
5. Hover over the snapshot to be backed up, and click its cloud backup icon ⊛ .
6. Select a cloud target.

Perform the following steps for each cloud target where the cloud backup will be stored. A new cloud backup job is created for each cloud target.

a. In the **Create Backup** dialog box, select a cloud target.

b. Select a format by clicking either **ZFS** or **Tar**.

   ZFS is the default value.

c. To select an incremental snapshot, click the check box for **Incremental**, and select a parent snapshot.

   Depending on whether the `Require parent exists` property is selected as described next, the parent snapshot must be available on both the local system and backed up to the same cloud target, or the parent snapshot does not also have to be available on the same cloud target. The parent snapshot must also be in the same format as the incremental snapshot: zfs or tar.

d. For an incremental snapshot and to require that it must have a parent snapshot on the same cloud target, select the check box for **Require parent exists**. For a parent snapshot and to require that you cannot later delete the cloud backup with the parent snapshot if it has cloud backed-up incremental snapshots (children), select the check box for **Require parent exists**.

   For both parent and incremental snapshots, this property can affect cloud backup restore and delete operations, as described in the following table.

Snapshot Type	Property Selected?	Restore Effect	Delete Effect
Parent	Yes	No effect.	Cannot delete cloud backup with parent snapshot if cloud backed-up incremental snapshots exist.
Parent	No	No effect.	Can delete cloud backup with parent snapshot if cloud backed-up incremental snapshots exist, which would save space on the cloud target.
Incremental	Yes	Parent snapshot must exist on the same cloud target, as well as on the local system.	No effect.
Incremental	No	Parent snapshot must exist on the local system. This allows you to create a cloud backup of an incremental snapshot without its parent snapshot on the same cloud target.	No effect.

e. Click **APPLY**.

   A **Backup Details** dialog box is displayed and dynamically shows the backup transfer rate, amount of data transferred, and status `in-progress`. The status changes to `completed` after completion. If the `write limit` property was set for the cloud target, the write traffic bandwidth is limited to the set value when uploading the cloud backup to the cloud target. Click **OK** to close the dialog box.

7. To view further progress details, including the completion percentage, from the **Configuration** menu, select **Services**, then click the **Jobs** tab.

   The **Jobs** list contains active and recently completed jobs in the cloud data service, and is in chronological order. The **Updates** column contains the date and time that the job was created, and the date and time that the backup was started/updated. The **Status** column indicates the completion percentage, the amount of data transferred, and the backup transfer rate.

8. To view job details, click on the job in the **Jobs** list.

   Click **OK** to close the dialog box. After the job has completed, the backup data will be available to other Oracle ZFS Storage Appliance systems configured with the same cloud target.

   If the tar format was selected, the cloud backup can be restored on any system, regardless of its operating system, that can access the cloud target.

Related Topics

- Canceling and Restarting a Cloud Job (BUI)
- Viewing Cloud Jobs (BUI)
- Defining a Cloud Target (BUI)
- Modifying a Cloud Target (BUI)
- Deleting a Cloud Target (BUI)
- Deleting a Cloud Backup (BUI)
- Cloud Properties

Creating a Cloud Backup (CLI)

After defining cloud targets, you can create a cloud backup of a share's (filesystem's or LUN's) full or incremental snapshot to a cloud target(s). Depending on how an incremental cloud backup is created, the parent snapshot must be available on both the local system and backed up to the same cloud target, or the parent snapshot does not also have to be available on the same cloud target. An incremental snapshot backup only backs up the difference between the previous snapshot and the current snapshot. Also, snapshot retention holds are preserved when moving the snapshot to another system via cloud snapshot backup.

Creating a cloud backup creates a cloud job with an operation of `backup`. Snapshots on different systems can have the same name, such as `snap2`, because snapshots are assigned a unique identification for each system.

There are two formats for cloud backups: "zfs" and "tar". The same snapshot on the same system can be used in two cloud backups: one with the zfs format, and one with the tar format. The formats support the features shown in the following table.

Feature	ZFS Format	Tar Format
Restore on only Oracle ZFS Storage Appliance or an Oracle Solaris server	✓	
Restore on any system regardless of operating system		✓
Supports both filesystem and LUN snapshots	✓	
Supports filesystem snapshots only		✓
Preserves filesystem properties	✓	

Feature	ZFS Format	Tar Format
Preserves LUN properties	✓	
Supports full and incremental backups	✓	✓
Supported within the same Oracle Cloud Infrastructure bucket	✓	✓
High-efficiency compression: • If the underlying share is compressed, less data is transferred and, therefore, the backup is faster. • Incremental backups are performed at the block level.	✓	
Mid-efficiency compression: • Files are read and compressed during the backup operation. • Incremental backups are performed at the file level. Especially not efficient if a large file is modified because the full file will be part of the incremental backup.		✓

> **Note:**
>
> After a local snapshot is in a cloud backup in your Oracle Cloud Infrastructure account, you can delete it locally, per your policies. However, retain local snapshots that could be parent snapshots for future incremental snapshots.

Filesystem snapshots that contain files created with the file retention feature (not the snapshot retention feature) are subject to the rules governing file retention. For information, see Planning Guidelines for File Retention.

Before You Begin

(Optional) To determine if a cloud backup of a snapshot is incremental and to view its parents, navigate to the snapshot (steps 1 through 7), enter the `targets` node, list and select the target, and list the target's properties. The `parents` property is populated if the cloud backup has parents, as shown in the following example.

```
hostname:shares default/fs-1@snap3> targets
hostname:shares default/fs-1@snap3 targets> select target-000
hostname:shares default/fs-1@snap3 target-000> ls
Properties:
           name = oci-phoenix
             id = a9aea6e0-55b8-4cd5-bdf0-ba0637be44b2
         format = zfs
        parents = snap1,snap2
```

1. Ensure the cloud target(s) is online.
   a. Go to `configuration services cloud targets`.
   b. Enter `ls` to list the cloud targets, and select the appropriate target.
   c. Enter `ls` and ensure that the `state` property is online. If it is offline, but the cloud data service is online, verify that the cloud target properties are valid. See Defining a Cloud Target (CLI).
2. Enter `top` to return to the root context.
3. Go to `shares` and enter `ls` to list the projects.

4. Select the project with the share snapshot to be backed up, and enter `ls` to list its properties and shares.

   ```
   hostname:shares> select default
   hostname:shares default> ls
   ```

5. Select the share with the snapshot to be backed up.

   ```
   hostname:shares default> select fs-1
   ```

6. Go to `snapshots` and enter `ls` to list the snapshots.

   ```
   hostname:shares default/fs-1> snapshots
   hostname:shares default/fs-1 snapshots> ls
   ```

7. Select the snapshot for cloud backup.

   ```
   hostname:shares default/fs-1 snapshots> select snap2
   hostname:shares default/fs-1@snap2>
   ```

8. Go to `backups`.

   ```
   hostname:shares default/fs-1@snap2> backups
   hostname:shares default/fs-1@snap2 backups>
   ```

9. Enter `create`.

   ```
   hostname:shares default/fs-1@snap2 backups> create
   hostname:shares default/fs-1@snap2 backup-001 (uncommitted)>
   ```

10. Set the target to an existing cloud target name.

    If the cloud target name has a space(s) in it, use command `select name="`*`target name`*`"`.

    Tip: Use tab completion for the cloud target name.

    ```
    hostname:shares default/fs-1@snap2 backup-001 (uncommitted)> set target=oci-phoenix
       target = oci-phoenix
    hostname:shares default/fs-1@snap2 backup-001 (uncommitted)>
    ```

11. To specify the tar format, set format to `tar`. Otherwise, the default `zfs` format is used.

    ```
    hostname:shares default/fs-1@snap2 backup-001 (uncommitted)> set format=tar
       format = tar
    hostname:shares default/fs-1@snap2 backup-001 (uncommitted)>
    ```

12. To set the current snapshot as an incremental snapshot, set `incremental` to `true`, and set the parent snapshot by name.

    Depending on how the `require_parent_exists` property is set as described next, the parent snapshot must be available on both the local system and backed up to the same cloud target, or the parent snapshot does not also have to be available on the same cloud target. The parent snapshot must also be in the same format as the incremental snapshot: zfs or tar.

    Tip: Use tab completion for a list of parent snapshot names on both the local system and on the cloud target.

    ```
    hostname:shares default/fs-1@snap2 backup-001 (uncommitted)> set incremental=true
       incremental = true
    hostname:shares default/fs-1@snap2 backup-001 (uncommitted)> set parent=snap0
       parent = snap0
    hostname:shares default/fs-1@snap2 backup-001 (uncommitted)>
    ```

13. For an incremental snapshot and to require that it must have a parent snapshot on the same cloud target, set `require_parent_exists` to `true`. For a parent snapshot and to require that you cannot later delete the cloud backup with the parent snapshot if it has cloud backed-up incremental snapshots (children), set `require_parent_exists` to `true`.

    For both parent and incremental snapshots, this property can affect cloud backup restore and delete operations, as described in the following table.

Snapshot Type	Property True?	Restore Effect	Delete Effect
Parent	Yes	No effect.	Cannot delete cloud backup with parent snapshot if cloud backed-up incremental snapshots exist.
Parent	No	No effect.	Can delete cloud backup with parent snapshot if cloud backed-up incremental snapshots exist, which would save space on the cloud target.
Incremental	Yes	Parent snapshot must exist on the same cloud target, as well as on the local system.	No effect.
Incremental	No	Parent snapshot must exist on the local system. This allows you to create a cloud backup of an incremental snapshot without its parent snapshot on the same cloud target.	No effect.

14. Enter `commit`.

    To send the cloud backup to a different cloud target, repeat steps 9 through 14 to create a new cloud backup job with a different cloud target.

```
hostname:shares default/fs-1@snap2 backup-001 (uncommitted)> commit
hostname:shares default/fs-1@snap2 backups>
```

15. To monitor progress as data is being backed up, enter `top` and go to `configuration services cloud jobs`.

```
hostname:configuration services cloud backups> top
hostname:> configuration services cloud jobs
```

16. Enter `ls` to view the cloud jobs.

    The `Jobs` list contains active and recently completed jobs in the cloud data service and is in chronological order. Up to one hundred cloud jobs are displayed. Use the `next` command to list the next 100 jobs, and use the `previous` command to list the previous 100 jobs.

```
hostname:configuration services cloud jobs> ls
Jobs:
JOB       OPERATION  CREATED
job-000   backup     2019-8-03 15:05:53
```

17. Find and select the cloud job by its number, and enter `ls` to view its details.

    After the job has completed, the backup data will be available to other Oracle ZFS Storage Appliance systems configured with the same cloud target. If the tar format was selected, the cloud backup can be restored on any system, regardless of its operating system, that can access the cloud target.

    The dataset path in the `details` property reflects the cloud backup format:

- ZFS format: `zfs/backups/zfs`
- Tar format: `zfs/backups/tar`

If the `write_limit` property was set for the cloud target, the write traffic bandwidth is limited to the set value when uploading the cloud backup to the cloud target.

```
hostname:configuration services cloud jobs> select job-000
hostname:configuration services cloud job-000> ls
Properties:
                    op = backup
                target = e641f83d-4628-42ba-8757-d66c4c98c0d9
            targetName = oci-phoenix
               created = 2020-10-22 22:23:40
               updated = 2020-10-22 22:27:33
                    id = 3babb944-07f6-4b69-8de1-f6dcfeab5fb2
                status = in-progress
                format = zfs
                  rate = 6MB/s
           transferred = 1.46G
        estimated_size = 40.1G
               dataset = p1/local/default/f-1
                backup = 3e035b7e546e0d02/1cbfdb5ff2259b76
              snapshot = snap9
               details = uploading backup to zfs/backups/zfs
                         3e035b7e546e0d02/1cbfdb5ff2259b76/000000001
```

Related Topics

- Canceling and Restarting a Cloud Job (CLI)
- Viewing Cloud Jobs (CLI)
- Defining a Cloud Target (CLI)
- Modifying a Cloud Target (CLI)
- Deleting a Cloud Target (CLI)
- Deleting a Cloud Backup (CLI)
- Cloud Properties

Viewing and Filtering Cloud Backups (BUI)

Use this procedure to view completed cloud backups on multiple pages and to filter them by using a basic search for either the cloud target name, source name, format, or dataset name. Use the advanced search feature to search for the same criteria simultaneously, plus you can search by the parent name.

1. From the **Configuration** menu, select **Services**, then **Cloud**.
2. Click the **Backups** tab.

   The **Backups** list contains completed cloud backups and is in chronological order.
3. If there are multiple pages of cloud backups, use the double arrow keys to navigate between the pages. To go to the first page, click the left-facing single arrow key. To go to the last page, click the right-facing single arrow key.
4. To filter the cloud backups using a basic search, click the search icon 🔍, enter the target name, source name, format, or dataset name, and press **Enter**.
   - You can only enter one property.
   - The target name must be an exact match; the other properties can be partial values.

- To view the entire, unfiltered list again, clear the search field and press **Enter**.

5. To filter the cloud backups using an advanced search, click the down-arrow icon ▾ below the search icon 🔍. In the **Search** dialog box, complete the fields as appropriate, and click **APPLY**.
   - Select the target and format from the pull-down lists. The source name, dataset name, and parent name can be partial values.
   - Upon completion, the search criteria and values appear in the search field beside the search icon.
   - To view the entire, unfiltered list again, clear the search field and press **Enter**.

Related Topics

- Restoring a Cloud Backup (BUI)
- Deleting a Cloud Backup (BUI)
- Cloud Properties

Viewing and Filtering Cloud Backups (CLI)

You can view 100 completed cloud backups at one time, and use the `next` and `previous` commands to scroll forward and backward by 100 backups. Additionally, you can filter the list by five properties: `target`, `source`, `format`, `dataset`, and `parent`.

1. Go to `configuration services cloud backups`.

   ```
   hostname:> configuration services cloud backups
   ```

2. Enter `ls` to view the properties and cloud backups.

   The `Backups` list contains completed cloud backups, and is in chronological order. Up to one hundred cloud backups are displayed.

   ```
   hostname:configuration services cloud backups> ls
   Properties:
      total    = 68
      selected = 68
      target   = (unset)
      format   = (unset)
      source   = (unset)
      dataset  = (unset)
      parent   = (unset)
   Backups:
   BACKUP     UPLOADED            TARGET      SOURCE   DATASET
   backup-000 2019-8-04 17:14:08 oci-phoenix hostname p1/local/default/fs-1@snap2
   backup-001 2019-8-04 17:25:36 oci-phoenix hostname p1/local/default/fs-1@snap31
   backup-002 2019-8-11 10:04:29 oci-phoenix hostname p1/local/default/fs-49@snap100
   backup-003 2019-8-15 11:18:32 oci-phoenix hostname p1/local/default/fs-1@snap1192
   backup-004 2019-8-16 23:09:16 oci-ashburn hostname p1/local/default/fs-22@snap107
   .
   .
   .
   backup-066 2019-8-19 11:28:47 oci-phoenix server38 p1/local/default/fs-386@snap51
   backup-067 2019-8-19 18:13:06 oci-phoenix server38 p1/local/default/fs-386@snap54
   ```

3. Use the `next` command to list the next 100 backups, and use the `previous` command to list the previous 100 backups.

4. To filter the list, `set` a property to a specific value, and then enter `ls`. To clear the filter, `unset` the property.

You can set and unset multiple filters. The properties that can be filtered are: `target`, `source`, `format`, `dataset`, and `parent`. The `target` and `parent` properties must be an exact match; the other properties can be partial values.

> **Note:**
>
> After filtering, the cloud backup numbering begins again at `backup-000`. You can immediately perform an operation on the new backup numbers. If you do not filter the cloud backups list and then perform an operation, be sure to use the original backup numbers.

```
hostname:configuration services cloud backups> set target=oci-ashburn
     target = oci-ashburn
hostname:configuration services cloud backups> ls
Properties:
   total    = 68
   selected = 1
   target   = oci-ashburn
   format   = (unset)
   source   = (unset)
   dataset  = (unset)
   parent   = (unset)
Backups:
BACKUP     UPLOADED           TARGET      SOURCE   DATASET
backup-000 2019-8-16 23:09:16 oci-ashburn hostname p1/local/default/fs-22@snap107
hostname:configuration services cloud backups> unset target
     target = (unset)
hostname:configuration services cloud backups> set source=server
     source = server
hostname:configuration services cloud backups> ls
Properties:
   total    = 68
   selected = 2
   target   = (unset)
   format   = (unset)
   source   = server
   dataset  = (unset)
   parent   = (unset)
Backups:
BACKUP     UPLOADED           TARGET      SOURCE   DATASET
backup-000 2019-8-19 11:28:47 oci-phoenix server38 p1/local/default/fs-386@snap51
backup-001 2019-8-19 18:13:06 oci-phoenix server38 p1/local/default/fs-386@snap54
hostname:configuration services cloud backups> unset source
     source = (unset)
hostname:configuration services cloud backups> set target=oci-phoenix
     target = oci-phoenix
hostname:configuration services cloud backups> set dataset=snap1
     dataset = snap1
hostname:configuration services cloud backups> ls
Properties:
   total    = 68
   selected = 2
   target   = oci-phoenix
   format   = (unset)
   source   = (unset)
   dataset  = snap1
   parent   = (unset)
Backups:
```

```
BACKUP     UPLOADED           TARGET      SOURCE   DATASET
backup-000 2019-8-11 10:04:29 oci-phoenix hostname p1/local/default/fs-49@snap100
backup-001 2019-8-15 11:18:32 oci-phoenix hostname p1/local/default/fs-1@snap1192
hostname:configuration services cloud backups> unset target
     target = (unset)
hostname:configuration services cloud backups> unset dataset
     dataset = (unset)
hostname:configuration services cloud backups> ls
Properties:
   total    = 68
   selected = 68
   target   = (unset)
   format   = (unset)
   source   = (unset)
   dataset  = (unset)
   parent   = (unset)
Backups:
BACKUP     UPLOADED           TARGET      SOURCE   DATASET
backup-000 2019-8-04 17:14:08 oci-phoenix hostname p1/local/default/fs-1@snap2
backup-001 2019-8-04 17:25:36 oci-phoenix hostname p1/local/default/fs-1@snap31
backup-002 2019-8-11 10:04:29 oci-phoenix hostname p1/local/default/fs-49@snap100
backup-003 2019-8-15 11:18:32 oci-phoenix hostname p1/local/default/fs-1@snap1192
backup-004 2019-8-16 23:09:16 oci-ashburn hostname p1/local/default/fs-22@snap107
.
.
.
backup-066 2019-8-19 11:28:47 oci-phoenix server38 p1/local/default/fs-386@snap51
backup-067 2019-8-19 18:13:06 oci-phoenix server38 p1/local/default/fs-386@snap54
```

Related Topics

- Restoring a Cloud Backup (CLI)
- Deleting a Cloud Backup (CLI)
- Cloud Properties

Restoring a Cloud Backup (BUI)

Cloud backups can be restored on any Oracle ZFS Storage Appliance system that has access to the cloud target. When a cloud backup with the zfs format contains a snapshot with a retention hold, that hold is preserved when the cloud backup is restored. Therefore, that snapshot and the share containing the snapshot cannot be deleted until the retention hold is released. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. For information about modifying a snapshot retention hold, see Editing a Snapshot Retention Policy (BUI).

Filesystem snapshots that contain files created with the file retention feature (not the snapshot retention feature) are subject to the rules governing file retention. For information, see Planning Guidelines for File Retention.

Filesystem cloud backups created with the tar format can be restored on Oracle ZFS Storage Appliance or on any system, regardless of its operating system, that can access the cloud target. After restoring a tar cloud backup on Oracle ZFS Storage Appliance, set the filesystem's properties, as described in Editing a Filesystem or LUN (BUI). If the filesystem snapshot had a retention hold, that property is not preserved. To set a retention hold, see Editing a Snapshot Retention Policy (BUI).

1. From the **Configuration** menu, select **Services**, then **Cloud**.
2. Click the **Backups** tab.

3. Find the cloud backup to be restored, and click its restore icon ↺ .
4. In the **Restore Backup** dialog box, select a pool and project. Either enter a new local share name to which to restore the dataset, or select **Use existing share** and enter an existing local share name.

   For a full (parent) cloud backup, if you enter a share name that has a parent snapshot in it, the parent snapshot is overwritten.

   For an incremental cloud backup, you can enter the same share name that has the parent snapshot. If you later repeat restoring an incremental cloud backup into the same share that has the parent snapshot, the old incremental snapshot is overwritten.
5. Click **APPLY**.

   After clicking **APPLY**, a **Restore Details** dialog box is displayed, and dynamically shows the restore transfer rate, amount of data transferred, and status `in-progress`. The status changes to `completed` after completion. If the `read limit` property was set for the cloud target, the read traffic bandwidth is limited to the set value when restoring the cloud backup from the cloud target. Click **OK** to close the dialog box.
6. To view further progress details, including the completion percentage, from the **Configuration** menu, select **Services**, then **Cloud**, and click the **Jobs** tab.

   The **Jobs** list contains active and recently completed jobs in the cloud data service, and is in chronological order. The **Updates** column contains the date and time that the job was created/updated, and the date and time that the restore was started. The **Status** column indicates the completion percentage, the amount of data transferred, and the restore transfer rate.
7. To view job details, click on the job in the **Jobs** list.

   Click **OK** to close the dialog box.
8. After the restore operation has completed, you can view cloud backup details:

   a. From the **Configuration** menu, select **Services**, then **Cloud**.

   b. Click the **Backups** tab.

      The **Backups** list contains completed cloud backups, and is in chronological order.

   c. Double-click on the cloud backup that was just restored.

   d. After viewing the **Backup Details** dialog box, click **OK**.

Related Topics

- Viewing and Filtering Cloud Backups (BUI)
- Deleting a Cloud Backup (BUI)
- Cloud Properties

Restoring a Cloud Backup (CLI)

Cloud backups can be restored on any Oracle ZFS Storage Appliance system that has access to the cloud target. When a cloud backup with the zfs format contains a snapshot with a retention hold, that hold is preserved when the cloud backup is restored. Therefore, that snapshot and the share containing the snapshot cannot be deleted until the retention hold is released. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. For information about modifying a snapshot retention hold, see Editing a Snapshot Retention Policy (CLI).

Filesystem snapshots that contain files created with the file retention feature (not the snapshot retention feature) are subject to the rules governing file retention. For information, see Planning Guidelines for File Retention.

Filesystem cloud backups created with the tar format can be restored on Oracle ZFS Storage Appliance or on any system, regardless of its operating system, that can access the cloud target. After restoring a tar cloud backup on Oracle ZFS Storage Appliance, set the filesystem's properties, as described in Editing a Filesystem or LUN (CLI). If the filesystem snapshot had a retention hold, that property is not preserved. To set a retention hold, see Editing a Snapshot Retention Policy (CLI).

1. Go to `configuration services cloud backups`.

   ```
   hostname:> configuration services cloud backups
   ```

2. Enter `ls` to view the properties and cloud backups.

   The `Backups` list contains completed cloud backups, and is in chronological order. Up to one hundred cloud backups are displayed at one time.

   ```
   hostname:configuration services cloud backups> ls
   Properties:
      total    = 68
      selected = 68
      target   = (unset)
      format   = (unset)
      source   = (unset)
      dataset  = (unset)
      parent   = (unset)
   Backups:
   BACKUP     UPLOADED           TARGET      SOURCE   DATASET
   backup-000 2019-8-04 17:14:08 oci-phoenix hostname p1/local/default/fs-1@snap2
   backup 001 2019-8-04 17:25:36 oci-phoenix hostname p1/local/default/fs-1@snap31
   .
   .
   .
   ```

3. Find and select the cloud backup to restore.

   ```
   hostname:configuration services cloud backups> select backup-000
   hostname:configuration services cloud backup-000>
   ```

4. Enter `restore`.

   ```
   hostname:configuration services cloud backups> restore
   hostname:configuration services cloud backup-000 restore (uncommitted)>
   ```

5. Enter `ls` to view the restore operation properties.

   ```
   hostname:configuration services cloud backup-000 restore (uncommitted)> ls
   Properties:
                  pool = (unset)
               project = (unset)
                 share = (unset)
   ```

6. For your system, set the pool, project, and share properties. Either enter a new local share name to which to restore the dataset, or set `useshare` to `true`, and enter an existing local share name. The default for `useshare` is `false`.

   For a full (parent) cloud backup, if you enter a share name that has a parent snapshot in it, the parent snapshot is overwritten.

For an incremental cloud backup, you can enter the same share name that has the parent snapshot. If you later repeat restoring an incremental cloud backup into the same share that has the parent snapshot, the old incremental snapshot is overwritten.

```
hostname:configuration services cloud backup-000 restore (uncommitted)> set pool=p1
hostname:configuration services cloud backup-000 restore (uncommitted)> set
project=default
hostname:configuration services cloud backup-000 restore (uncommitted)> set
useshare=true
hostname:configuration services cloud backup-000 restore (uncommitted)> set
share=existing-share-name
hostname:configuration services cloud backup-000 restore (uncommitted)>
```

7. Enter `commit`.

```
hostname:configuration services cloud backup-000 restore (uncommitted)> commit
hostname:configuration services cloud>
```

8. Go to `jobs` and enter `ls` to view the cloud jobs.

   The `Jobs` list contains active and recently completed jobs in the cloud data service, and is in chronological order.

```
hostname:configuration services cloud jobs> ls
Jobs:
JOB       OPERATION  CREATED
job-000   backup     2019-8-03 15:05:53
job-001   restore    2019-8-06 17:14:36
```

9. Select the cloud job by its number, and enter `ls` to view its details.

   The restored cloud backup will not be available from configured storage protocols until the restore job has completed.

   If the `write_limit` property was set for the cloud target, the write traffic bandwidth is limited to the set value when uploading the cloud backup to the cloud target.

```
hostname:configuration services cloud jobs> select job-001
hostname:configuration services cloud job-001> ls
                     op = restore
                 target = 9ca87404-24e2-11e9-a929-0b69dceb9c81
             targetName = oci-phoenix
                created = 2019-8-06 17:14:36
                updated = 2019-8-11 20:33:09
                     id = cbb88888852e58088/445c15865cb329c3
                 status = in-progress
                 format = zfs
                   rate = 74KB/s
            transferred = 5.98G
         estimated_size = 1.00T
                dataset = p1/local/default/fs-1
                 backup = 3e035b7e546e0d02/1cbfdb5ff2259b76
               snapshot = snap2
                details = uploading backup to zfs/backups/zfs/3e035b7e546e0d02/
db7fd6c55558cea0/000000001
```

10. After the restore operation has completed, you can view cloud backup details:

    a. Go to `configuration services cloud backups`.

    b. Select the cloud backup by its number.

    ```
    hostname:configuration services cloud backups> select backup-000
    ```

    c. Enter `ls` to view the cloud backup details.

```
hostname:configuration services cloud backup-000> ls
Properties:
              target = oci-phoenix
              source = hostname
             dataset = p1/local/default/fs-1@snap2
              format = zfs
                tier = standard
                size = 1.00T
             started = 2019-8-03 15:05:53
            uploaded = 2019-8-04 17:14:08
                  id = cbb88888852e58088/445c15865cb329c3
```

Related Topics

- Viewing and Filtering Cloud Backups (CLI)
- Deleting a Cloud Backup (CLI)
- Cloud Properties

Managing the Cloud Service

To manage cloud storage, use these tasks:

- Deleting a Cloud Backup - BUI, CLI
- Viewing Cloud Jobs - BUI, CLI
- Canceling and Restarting a Cloud Job - BUI, CLI
- Modifying a Cloud Target - BUI, CLI
- Deleting a Cloud Target - BUI, CLI

Deleting a Cloud Backup (BUI)

After a cloud backup on Oracle ZFS Storage Appliance is no longer needed, delete it using the appliance. Note that after a cloud backup delete operation is started, it cannot be cancelled.

If, on any appliance, a cloud backup restore operation is in progress that includes the cloud backup data, you cannot delete that cloud backup unless you confirm your action.

If the cloud backup is a parent snapshot with cloud backed-up incremental snapshots (children), you can delete the parent cloud backup if the incremental cloud backups were created without the property `Require parent exists` selected (on the cloud target). If the property `Require parent exists` was selected for any of the incremental snapshots, you cannot delete the parent snapshot from the cloud target. For more information, see Creating a Cloud Backup (BUI).

If the cloud backup contains a share with a snapshot with a retention hold, that snapshot and its share cannot be deleted until the retention hold is released. Therefore, the cloud backup cannot be deleted. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. For information about modifying a snapshot retention hold, see Editing a Snapshot Retention Policy (BUI).

Filesystem snapshots that contain files created with the file retention feature (not the snapshot retention feature) are subject to the rules governing file retention. For information, see Planning Guidelines for File Retention.

1. From the **Configuration** menu, select **Services**, then **Cloud**.
2. Click the **Backups** tab.

The **Backups** list contains completed cloud backups, and is in chronological order.

3. Find and hover over the cloud backup to be deleted, and click its delete icon 🗑 .

   You must force the delete action if the cloud backup is a parent snapshot with cloud backed-up incremental snapshots (children) that were created without the property `Require parent exists` selected (on the cloud target). To do so, hold the Shift key while clicking the delete icon.

4. Confirm your action by clicking **OK**.

Related Topics

- Viewing and Filtering Cloud Backups (BUI)
- Restoring a Cloud Backup (BUI)
- Cloud Properties

Deleting a Cloud Backup (CLI)

After a cloud backup on Oracle ZFS Storage Appliance is no longer needed, delete it using the appliance. Note that after a cloud backup delete operation is started, it cannot be cancelled.

If, on any appliance, a cloud backup restore operation is in progress that includes the cloud backup data, you cannot delete that cloud backup unless you confirm your action.

If the cloud backup is a parent snapshot with cloud backed-up incremental snapshots (children), you can delete the parent cloud backup if the incremental cloud backups were created with the property `require_parent_exists` (on the cloud target) set to `false`. If the property `require_parent_exists` was set to `true` for any of the incremental snapshots, you cannot delete the parent snapshot from the cloud target. For more information, see Creating a Cloud Backup (CLI).

If the cloud backup contains a share with a snapshot with a retention hold, that snapshot and its share cannot be deleted until the retention hold is released. Therefore, the cloud backup cannot be deleted. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. For information about modifying a snapshot retention hold, see Editing a Snapshot Retention Policy (CLI).

Filesystem snapshots that contain files created with the file retention feature (not the snapshot retention feature) are subject to the rules governing file retention. For information, see Planning Guidelines for File Retention.

1. Go to `configuration services cloud backups`.

   ```
   hostname:> configuration services cloud backups
   ```

2. Enter `ls` to view the properties and cloud backups.

   The `Backups` list contains completed cloud backups, and is in chronological order.

   ```
   hostname:configuration services cloud backups> ls
   Properties:
      total    = 68
      selected = 68
      target   = (unset)
      format   = (unset)
      source   = (unset)
      dataset  = (unset)
   Backups:
   BACKUP     UPLOADED           TARGET      SOURCE   DATASET
   backup-000 2019-8-04 17:14:08 oci-phoenix hostname p1/local/default/fs-1@snap2
   ```

```
backup 001 2019-8-04 17:25:36 oci-phoenix hostname p1/local/default/fs-1@snap31
.
.
.
```

3. Find the cloud backup to delete.
4. Enter `destroy` and the cloud backup name.

   ```
   hostname:configuration services cloud backups> destroy backup-000
   This will destroy "backup-000". Are you sure? (Y/N)
   ```

   You must force the delete action if the cloud backup is a parent snapshot with cloud backed-up incremental snapshots (children) that were created with the property `require_parent_exists` (on the cloud target) set to `false`. To do so, enter option `-f` before the cloud backup name.

   ```
   hostname:configuration services cloud backups> destroy -f backup-000
   This will destroy "backup-000". Are you sure? (Y/N)
   ```

5. Confirm your action by entering `Y`.

Related Topics

- Viewing and Filtering Cloud Backups (CLI)
- Restoring a Cloud Backup (CLI)
- Cloud Properties

Viewing Cloud Jobs (BUI)

The **Jobs** list contains active and recently completed jobs in the cloud data service, and is in chronological order. You can filter the **Jobs** list by target name, operation, and status.

1. From the **Configuration** menu, select **Services**, then **Cloud**.
2. Click the **Jobs** tab.
   - **Icon to left of job** - A green icon indicates a successful job, and a gray icon indicates a cancelled job or an error.
   - **Target**: Cloud target name.
   - **Updates** - Date and time that the job was created, and the date and time that the operation was started/updated.
   - **Operation** - Indicates an operation of **backup**, **restore**, or **delete**.
   - **Status** - Indicates the completion percentage, the amount of data transferred, and the transfer rate. If a backup operation and the `write limit` property was set for the cloud target, the write traffic bandwidth is limited to the set value. If a restore operation and the `read limit` property was set for the cloud target, the read traffic bandwidth is limited to the set value.
3. To view job details, find and click on a job in the **Jobs** list.
   - If there are multiple pages of cloud jobs, use the double arrow keys to navigate between the pages. To go to the first page, click the left-facing single arrow key. To go to the last page, click the right-facing single arrow key.
   - To search for a cloud job by target name, operation, and status, click the down-arrow icon ▾ below the search icon 🔍. In the **Search** dialog box, complete the fields as appropriate, and click **APPLY**.

Upon completion, the search criteria and values appear in the search field beside the search icon. To view the entire, unfiltered list again, clear the search field and press **Enter**.

- For a description of each detailed cloud job property, see Cloud Properties.
- When finished viewing the job details, click **OK** to close the **Job Details** dialog box.

Related Topics

Cloud Properties

Viewing Cloud Jobs (CLI)

The `Jobs` list contains active and recently completed jobs in the cloud data service, and is in chronological order.

1. Go to `configuration services cloud jobs`.

   ```
   hostname:configuration services cloud jobs>
   ```

2. Enter `ls` to view the cloud jobs.

   - `JOB` - Cloud job number.
   - `OPERATION` - Indicates an operation of `backup`, `restore`, or `delete`.
   - `CREATED` - Date and time that the cloud job was created.

   ```
   hostname:configuration services cloud jobs> ls
   Jobs:
   JOB       OPERATION  CREATED
   job-000   backup     2019-8-03 15:05:53
   ```

3. Find and select the cloud job by its number, and enter `ls` to view its details.

   `status` values: `in-progress`, `pending`, and `completed`. If an `in-progress` backup operation and the `write_limit` property was set for the cloud target, the write traffic bandwidth is limited to the set value. If an `in-progress` restore operation and the `read_limit` property was set for the cloud target, the read traffic bandwidth is limited to the set value. For a description of all detailed cloud job properties, see Cloud Properties.

   ```
   hostname:configuration services cloud jobs> select job-000
   hostname:configuration services cloud job-000> ls
   Properties:
                       op = backup
                   target = e641f83d-4628-42ba-8757-d66c4c98c0d9
               targetName = oci-phoenix
                  created = 2019-8-03 15:05:53
                  updated = 2019-8-03 22:27:33
                       id = 3babb944-07f6-4b69-8de1-f6dcfeab5fb2
                   status = in-progress
                   format = zfs
                     rate = 6MB/s
              transferred = 1.46G
           estimated_size = 40.1G
                  dataset = p1/local/default/f-1
                   backup = 3e035b7e546e0d02/1cbfdb5ff2259b76
                 snapshot = snap9
                  details = uploading backup to zfs/backups/
                            3e035b7e546e0d02/1cbfdb5ff2259b76/000000001
   ```

Related Topics

Cloud Properties

Canceling and Restarting a Cloud Job (BUI)

Canceling a cloud job stops the job, but does not immediately delete it from the **Jobs** list on Oracle ZFS Storage Appliance. Restarting a cloud job restarts the job from the beginning. Separately, note that you cannot cancel a cloud backup delete operation after it has started.

1. From the **Configuration** menu, select **Services**, then **Cloud**.
2. Click the **Jobs** tab.

   The list contains active and recently completed jobs in the cloud data service, and is in chronological order.
3. Find and hover over the job to cancel, and click its cancel icon ⊗ .
4. Confirm your cancellation action by clicking **OK**.
5. To restart a cloud job, find and hover over the job to restart, and click its restart icon ↻ .

Related Topics

- Creating a Cloud Backup (BUI)
- Cloud Properties

Canceling and Restarting a Cloud Job (CLI)

Canceling a cloud job stops the job, but does not immediately delete it from the `Jobs` list. Restarting a cloud job restarts the job from the beginning. Separately, note that you cannot cancel a cloud backup delete operation after it has started.

1. Go to `configuration services cloud jobs`.

   ```
   hostname:configuration services cloud jobs>
   ```
2. Enter `ls` to view the cloud jobs.

   The `Jobs` list contains active and recently completed jobs in the cloud data service, and is in chronological order.

   ```
   hostname:configuration services cloud jobs> ls
   Jobs:
   JOB       OPERATION  CREATED
   job-000   backup     2019-8-03 15:05:53
   ```
3. Find and select the cloud job to cancel.

   ```
   hostname:configuration services cloud jobs> select job-000
   ```
4. Enter `cancel`.

   ```
   hostname:configuration services cloud job-000> cancel
   ```
5. To restart a cloud job, find and select the job to restart, and enter `restart`.

   ```
   hostname:configuration services cloud job-000> restart
   ```

Related Topics

- Creating a Cloud Backup (CLI)
- Cloud Properties

Modifying a Cloud Target (BUI)

A cloud target's name and web proxy settings can be modified. Optionally, the write and read limit bandwidths can also be changed.

1. From the **Configuration** menu, select **Services**, then **Cloud**.
2. Click the **Targets** tab.
3. Hover over the cloud target to modify, and click its edit icon ✎ .
4. Enter values for the desired fields, and click **APPLY**.
   - **Name** - Your name for this cloud target, which must be unique for your system.
   - **Use web proxy** - Select this check box to use a proxy for system communications with the web, and complete the fields for host port name and number. The user name and password fields are optional.
   - **Write limit bandwidth** (optional) - Select this check box to limit the traffic write bandwidth when uploading a cloud backup to the cloud target. Enter a value and select a unit of measurement. For example, `5 M/s` limits writes to the cloud target to 5 megabytes per second.
   - **Read limit bandwidth** (optional) - Select this check box to limit the traffic read bandwidth when restoring a cloud backup from the cloud target. Enter a value and select a unit of measurement. For example, `4 M/s` limits reads from the cloud target to 4 megabytes per second.

Related Topics

- Defining a Cloud Target (BUI)
- Deleting a Cloud Target (BUI)
- Cloud Properties

Modifying a Cloud Target (CLI)

A cloud target's name and web proxy settings can be modified. Optionally, the write and read limit bandwidths can also be changed.

1. Go to `configuration services cloud targets`.

```
hostname:configuration services cloud targets>
```

2. Find and select the cloud target to modify, and enter `get` to view its properties.

```
hostname:configuration services cloud target-000> get
Properties:
                      name = oci-phoenix
                  location = https://objectstorage.us-phoenix-1.oraclecloud.com
                      user = ocid1.user.oc1..aaa56chx6rcm53g4tij74pymqffm4gsxlrhnq
                    bucket = test
                   tenancy = ocid1.tenancy.oc1..aaaaao6lmlzrvmk2x3uv7cglgxpan5ldsmq
                       key = true
                  proxy_on = true
                proxy_host = www-proxy.us.example.com:80
                proxy_user =
            proxy_password =
                    online = true
                        id = 77e3a201-b4e0-4a2c-9f9d-aee21eda9954
```

3. To modify the properties, use the `set` command, such as `set name=oci-phoenix-test`. The following properties can be modified:
    - `name` - Your name for this cloud target, which must be unique for your system.
    - `proxy_on` - Set this to `true` to use a proxy for system communications with the web, and complete the property for `proxy_host` (host port name and number). The `proxy_user` (user name) and `proxy_password` (user password) properties are optional. Example for `proxy_host`: `www-proxy.us.example.com:80`.
    - `writelimit` (optional) - Set to a value and unit of measurement to limit the traffic write bandwidth when uploading a cloud backup to the cloud target. For example, `5M/s` limits writes to the cloud target to 5 megabytes per second.
    - `readlimit` (optional) - Set to a value and unit of measurement to limit the traffic read bandwidth when restoring a cloud backup from the cloud target. For example, `4M/s` limits reads from the cloud target to 4 megabytes per second.
4. When done, enter `commit`.

```
hostname:configuration services cloud target-000 (uncommitted)> commit
```

Related Topics

- Defining a Cloud Target (CLI)
- Deleting a Cloud Target (CLI)
- Cloud Properties

Deleting a Cloud Target (BUI)

After a cloud target is no longer needed, delete it using Oracle ZFS Storage Appliance. Deleting a cloud target does not destroy its bucket nor the objects in your Oracle Cloud Infrastructure account. Note that after a cloud target delete operation is started, it cannot be canceled.

If a snapshot is configured to use the cloud target, a warning is displayed. You can confirm your action to delete the cloud target or you can cancel the action and abort jobs submitted to the cloud target. For informational reasons, snapshots that have already been backed up will show the deleted cloud target as the backup target.

If a new cloud target is later created that points to the same location, it will have a new identification. Therefore, any cloud actions, such as backup and restore operations, will need to be recreated for the new cloud target.

1. From the **Configuration** menu, select **Services**, then **Cloud**.
2. Click the **Targets** tab.
3. Hover over the cloud target to delete, and click its trash icon 🗑 .
4. Confirm your action by clicking **OK**.

Related Topics

- Modifying a Cloud Target (BUI)
- Cloud Properties

Deleting a Cloud Target (CLI)

After a cloud target is no longer needed, delete it using Oracle ZFS Storage Appliance. Deleting a cloud target does not destroy its bucket nor the objects in your Oracle Cloud Infrastructure account. Note that after a cloud target delete operation is started, it cannot be canceled.

If a snapshot is configured to use the cloud target, a warning is displayed. You can confirm your action to delete the cloud target or you can cancel the action and abort jobs submitted to the cloud target. For informational reasons, snapshots that have already been backed up will show the deleted cloud target as the backup target.

If a new cloud target is later created that points to the same location, it will have a new identification. Therefore, any cloud actions, such as backup and restore operations, will need to be recreated for the new cloud target.

1. Go to `configuration services cloud targets`.

   ```
   hostname:configuration services cloud targets>
   ```

2. Find the cloud target to delete.
3. Enter `destroy` and the cloud target name.

   ```
   hostname:configuration services cloud targets> destroy target-000
   This will destroy "target-000". Are you sure? (Y/N)
   ```

4. Confirm your action by entering `Y`.

Related Topics

- Modifying a Cloud Target (CLI)
- Cloud Properties

Cloud Properties

The cloud properties are grouped into the following categories:

- Cloud Data Service Properties
- Cloud Target Properties
- Cloud Job Properties
- Cloud Backup Properties
- Cloud Backup Restore Properties

Cloud data service properties are used when configuring the cloud data service.

Table 2-12 Cloud Data Service Properties

BUI Field Name	CLI Name	Description
green status icon	`status = online`	Indicates that the cloud data service is online.
gray status icon	`status = offline`	Indicates that the cloud data service is offline.

Table 2-12 (Cont.) Cloud Data Service Properties

BUI Field Name	CLI Name	Description
SSL/TLS versions and ciphers	`tls_version`	For the BUI, select or clear check boxes for SSL/TLS versions. For the CLI, set the property to specific values, separated by commas. Use tab completion to see a list of possible values.
SSL/TLS versions and ciphers	`ciphers`	For the BUI, select or clear check boxes for ciphers. For the CLI, set the property to specific values, separated by commas. Use tab completion to see a list of possible values.

Cloud target properties are used when defining or modifying a cloud target.

Table 2-13 Cloud Target Properties

BUI Field Name	CLI Name	Description
Name	`name`	Cloud target name; must be unique to the system. This property can be modified after the cloud target is defined.
Location	`location`	Cloud target location, which is a regional API endpoint URL. See also Setting Up Your Oracle Cloud Infrastructure Account.
User	`user`	User name Oracle Cloud Identifier (OCID). See also Setting Up Your Oracle Cloud Infrastructure Account.
Use data bucket	`databucket_on`	Indicates if your data and metadata are in different buckets. Data will be saved to the `Data Bucket`, and metadata will be saved to the `Bucket`. If you use an archive bucket, the metadata must be saved to a standard bucket. For the BUI, select or clear the associated check box. For the CLI, set the property to `true` or `false`.
Bucket	`bucket`	An established bucket name. The combination of `location`, `bucket`, and `data_bucket` must be unique for your system. If in the BUI you selected `Use data bucket` or if in the CLI you set `databucket_on` to `true`, this property is for the metadata bucket. If in the BUI you did not select `Use data bucket` or if in the CLI you set `databucket_on` to `false`, this property is for both the metadata and data.
Data Bucket	`data_bucket`	An established bucket name. The combination of `location`, `bucket`, and `data_bucket` must be unique for your system. If in the BUI you selected `Use data bucket` or if in the CLI you set `databucket_on` to `true`, this property is for the data bucket. If in the BUI you did not select `Use data bucket` or if in the CLI you set `databucket_on` to `false`, this property cannot be set.
Tenancy	`tenancy`	Tenancy OCID. See also Setting Up Your Oracle Cloud Infrastructure Account.
Private Key	`key`	Appliance private key. See also Setting Up Your Oracle Cloud Infrastructure Account.

Table 2-13 (Cont.) Cloud Target Properties

BUI Field Name	CLI Name	Description
Use web proxy	`proxy_on`	Indicates if a proxy is used for system communications with the web. For the BUI, select or clear the associated check box. For the CLI, set the property to `true` or `false`. This property can be modified after the cloud target is defined.
Host : port	`proxy_host`	Host name and port number for the web proxy. This property can be modified after the cloud target is defined.
Username	`proxy_user`	User name needed to access the web proxy. This property can be modified after the cloud target is defined.
Password	`proxy_password`	Password needed to access the web proxy. This property can be modified after the cloud target is defined.
green status icon	`online = true`	Indicates that the cloud target is online and available.
Write limit bandwidth	`writelimit`	Value and unit of measurement for limiting the traffic write bandwidth when uploading a cloud backup to the cloud target. Units of measure: • `B/s` for bytes per second • `K/s` for kilobytes per second • `M/s` for megabytes per second • `G/s` for gigabytes per second
Read limit bandwidth	`readlimit`	Value and unit of measurement for limiting the traffic read bandwidth when restoring a cloud backup from the cloud target. Units of measure: • `B/s` for bytes per second • `K/s` for kilobytes per second • `M/s` for megabytes per second • `G/s` for gigabytes per second
gray status icon	`online = false`	Indicates that the cloud target is offline and unavailable.
	`id`	Cloud target identification.

Cloud job properties are used when viewing a cloud job.

Table 2-14 Cloud Job Properties

BUI Field Name	CLI Name	Description
Operation	`op`	Operation type: `backup`, `delete`, or `restore`.
Target	`target`	Cloud target identification.
	`targetName`	Cloud target name.
Updates	`created`	Creation date and time of cloud job.
	`updated`	Update/start date and time of cloud job.
	`id`	Cloud job identification.
Format	`format`	Cloud backup format: `zfs` (default) or `tar`.
Status	`status`	Cloud job status: `in-progress`, `pending`, `completed`, or `error`.

Table 2-14 (Cont.) Cloud Job Properties

BUI Field Name	CLI Name	Description
	`rate`	Rate of data transfer, expressed in a unit of measure (see "Amount of data transferred") per second.
	`transferred` and `estimated_size`	Amount of data transferred and estimated size. If job is in progress, the BUI displays the percentage of data transferred and the estimated size; in the CLI, the `transferred` and `estimated_size` properties can be used to calculate the percentage. Units of measure: • `B` for bytes • `K` for kilobytes • `M` for megabytes • `G` for gigabytes • `T` for terabytes • `P` for petabytes • `E` for exabytes
	`dataset`	Dataset. Defined as *`pool name`*`/local/`*`project name`*`/`*`share name`*.
	`backup`	Cloud backup identification. Example: `3e035b7e546e0d02 / db7fd6c55558cea0`
	`snapshot`	Snapshot name.
	`details`	Cloud job details, including the status of the cloud backup, the bucket name, the cloud backup format, and the backup identification. Example: `uploading backup to zfs / backups / zfs / 3e035b7e546e0d02 / db7fd6c55558cea0 / 000000001`

Cloud backup properties are used when creating or viewing a cloud backup.

Table 2-15 Cloud Backup Properties

BUI Field Name	CLI Name	Description
Target	`target`	Cloud target name.
Source	`source`	Cloud backup source system.
Dataset	`dataset`	Dataset. Defined as *`pool name`*`/local/`*`project name`*`/`*`share name`*`@`*`snapshot name`*.
Format	`format`	Cloud backup format: `zfs` (default) or `tar`.
Storage Tier	`tier`	Storage tier type: `standard` or `archive`.
Size	`size`	Size of the cloud backup. Units of measure: • `B` for bytes • `K` for kilobytes • `M` for megabytes • `G` for gigabytes • `T` for terabytes • `P` for petabytes • `E` for exabytes

Table 2-15 (Cont.) Cloud Backup Properties

BUI Field Name	CLI Name	Description
Started	`started`	Start date and time of cloud backup.
Uploaded	`uploaded`	Uploaded date and time of cloud backup.
	`id`	Cloud backup identification. Example: `3e035b7e546e0d02 / db7fd6c55558cea0`
Parent	`parent`	Parent snapshot for an incremental snapshot. This property is displayed only when an incremental snapshot is selected.
Require parent exists	`require_parent_exists`	For an incremental snapshot, indicates if the parent snapshot must exist on the same cloud target, as well as on the local system. For the BUI, select or clear the associated check box. For the CLI, set the property to `true` or `false`.
Incremental	`incremental`	Indicates if this is an incremental snapshot. For the BUI, select or clear the associated check box. For the CLI, set the property to `true` or `false`.

Cloud backup restore properties are used when restoring a cloud backup on Oracle ZFS Storage Appliance.

After restoring a tar cloud backup on Oracle ZFS Storage Appliance, set the filesystem's properties because this type of backup does not preserve the original filesystem properties nor does it inherit project properties.

Table 2-16 Cloud Backup Restore Properties

BUI Field Name	CLI Name	Description
Pool	`pool`	Local pool name for the cloud backup restore.
Project	`project`	Local project name for the cloud backup restore.
Share	`share`	Local share name for the cloud backup restore.
Use existing share	`useshare`	Indicates if the cloud backup will be restored into an existing local share. For the BUI, select or clear the associated check box. For the CLI, set the property to `true` or `false`.

Understanding the Appliance Status

The Status section provides a summary of Oracle ZFS Storage Appliance status and configuration options. Use the following sections for conceptual and procedural information about appliance status views and related service configuration:

- About Oracle ZFS Storage Appliance
- Status Dashboard
- Summary of Pool Usage
- Summary of Memory Usage
- Disk Activity Dashboard
- Dashboard CLI
- Running the Dashboard Continuously

- Status Dashboard Settings
- Changing the Displayed Activity Statistics
- Changing the Activity Thresholds
- NDMP Status
- NDMP States

Status Dashboard

The dashboard summarizes appliance status.

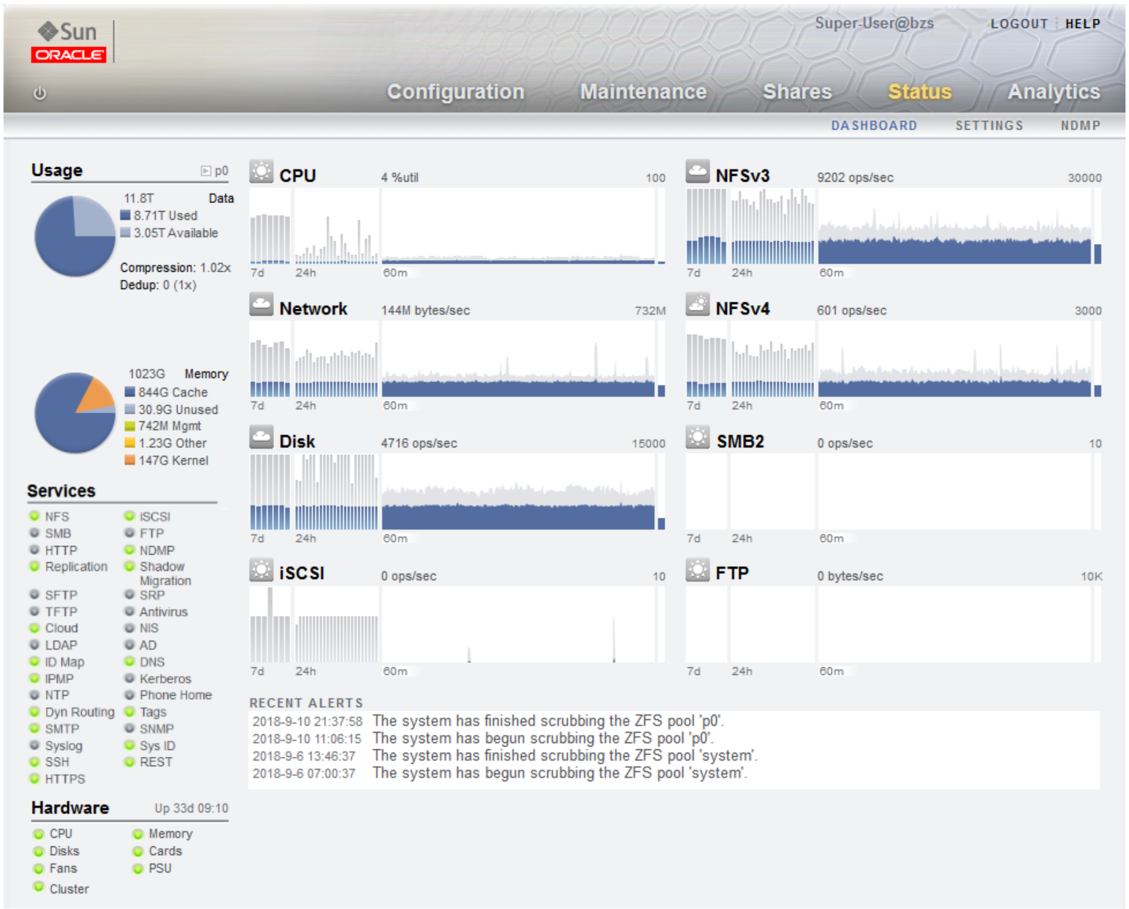

2018-9-10 21:37:58 The system has finished scrubbing the ZFS pool 'p0'.
2018-9-10 11:06:15 The system has begun scrubbing the ZFS pool 'p0'.
2018-9-6 13:46:37 The system has finished scrubbing the ZFS pool 'system'.
2018-9-6 07:00:37 The system has begun scrubbing the ZFS pool 'system'.

The status dashboard provides links to all main screens of the browser user interface (BUI). Over 100 visible items on the dashboard link to associated BUI screens indicated by a border or highlighted text that appears when hovered over. The sections that follow describe the areas of the dashboard in detail.

Usage Dashboard

The **Usage** area of the dashboard provides a summary of your storage pool and main memory usage. The name of the pool appears at the top right of the **Usage** area. If multiple pools are configured, use the pull-down list to select the pool you want to display.

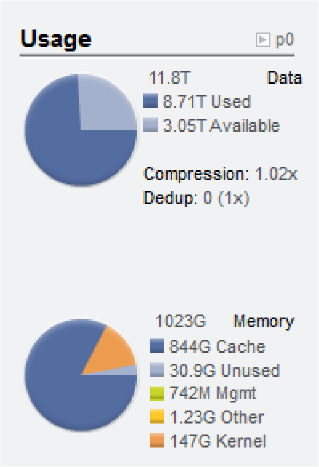

Usage ▷ p0

11.8T Data
- ■ 8.71T Used
- ■ 3.05T Available

Compression: 1.02x
Dedup: 0 (1x)

1023G Memory
- ■ 844G Cache
- ■ 30.9G Unused
- ■ 742M Mgmt
- ■ 1.23G Other
- ■ 147G Kernel

The total pool capacity is displayed to the right of the storage usage pie chart. The storage pie chart details the used and available space. To go to the **Shares** screen for the pool, click the storage pie chart.

The total system physical memory is displayed to the right of the memory pie chart. To the left is a pie chart showing memory usage by component. To go to the Analytics worksheet for dynamic memory usage broken down by application name, click the **Memory** pie chart.

Services Dashboard

The **Services** area of the dashboard shows the status of services on the appliance, with a light icon to show the state of each service.

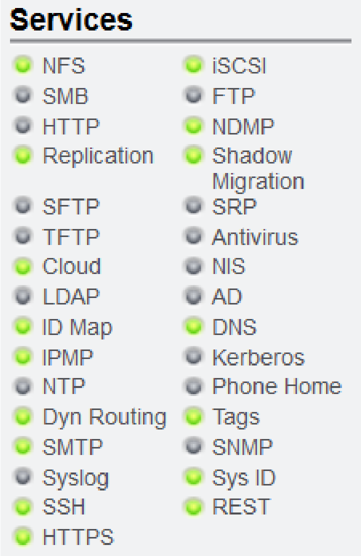

Services

🟢 NFS
⚪ SMB
⚪ HTTP
🟢 Replication
⚪ SFTP
⚪ TFTP
🟢 Cloud
⚪ LDAP
🟢 ID Map
🟢 IPMP
⚪ NTP
🟢 Dyn Routing
🟢 SMTP
⚪ Syslog
🟢 SSH
🟢 HTTPS

🟢 iSCSI
⚪ FTP
🟢 NDMP
🟢 Shadow Migration
⚪ SRP
⚪ Antivirus
⚪ NIS
⚪ AD
🟢 DNS
⚪ Kerberos
⚪ Phone Home
🟢 Tags
⚪ SNMP
🟢 Sys ID
🟢 REST

Most services are green to indicate that the service is online, or grey to indicate that the service is disabled. For a reference of all possible states and icon colors, see Browser User Interface (BUI).

To go to the associated **Configuration** screen, click on a service name. The **Properties** screen appears with configurable fields, restart, enable, and disable icons, and a link to the associated **Logs** screen for the service.

Hardware Dashboard

The **Hardware** area of the dashboard shows an overview of hardware on the appliance.

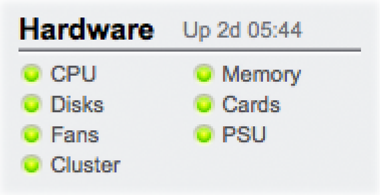

Hardware Up 2d 05:44

● CPU	● Memory
● Disks	● Cards
● Fans	● PSU
● Cluster	

If there is a known fault, the amber fault ● icon appears.

To go to the **Hardware** screen for a detailed look at hardware state, click the name of a hardware component.

Activity Dashboard

The activity area of the dashboard shows graphs of eight performance statistics by default. The example in this section shows **Disk** operations/second. The statistical average is plotted in blue and the maximum appears in light gray.

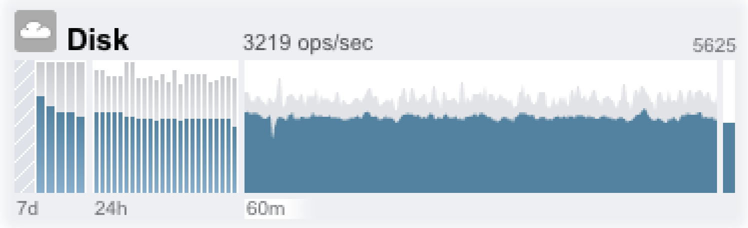

To go to the **Analytics** worksheet for an activity, click one of the four graphs (day, hour, minute, second) for the statistic you want to evaluate.

To view the average for each graph, mouse-over a graph, and the average appears in the tooltip. The weather icon in the upper-left provides a report of activity according to thresholds you can customize for each statistic on the Status Dashboard Settings screen.

Table 2-17 Summary of Statistic Graphs

Statistic Graph	Description
7-day graph (7d)	A bar chart, with each bar representing one day.
24-hour graph (24h)	A bar chart, with each bar representing one hour.
60-minute graph (60m)	A line plot, representing activity over one hour (also visible as the first one-hour bar in the 24-hour graph).
1-second graph	A line plot, representing instantaneous activity reporting.

The average for the selected plot is shown numerically above the graph. To change the average that appears, select the average you want, either 7d, 24h, or 60m.

The vertical scale of all graphs is printed on the top right, and all graphs are scaled to this same height. The height is calculated from the selected graph (plus a margin). The height will rescale based on activity in the selected graph, with the exception of utilization graphs which have a fixed height of 100 percent.

Since the height can rescale, 60 minutes of idle activity may look similar to 60 minutes of busy activity. Always check the height of the graphs before trying to interpret what they mean.

Understanding some statistics may not be obvious - you might wonder, for a particular appliance in your environment, whether 1000 NFSv3 ops/sec is considered busy or idle. This is where the 24-hour and 7-day plots can help, to provide historic data next to the current activity for comparison.

The plot height is calculated from the selected plot. By default, the 60-minute plot is selected. So, the height is the maximum activity during that 60-minute interval (plus a margin). To rescale all plots to span the highest activity during the previous 7 days, select 7d. This makes it easy to see how current activity compares to the last day or week.

The weather icon is intended to grab your attention when something is unusually busy or idle. To go to the weather threshold configuration page, click the weather icon. There is no good or bad threshold, rather the BUI provides a gradient of levels for each activity statistic. The statistics on which weather icons are based provide an *approximate* understanding for appliance performance that you should customize to your workload, as follows:

- Different environments have different acceptable levels for performance (latency), and so there is no one-size-fits-all threshold.
- The statistics on the **Dashboard** are based on operations/sec and bytes/sec, so you should use **Analytics** worksheets for an accurate understanding of system performance.

Recent Alerts

This section shows the last four appliance alerts. Click the box to go to the **Logs** screen to examine all recent alerts in detail.

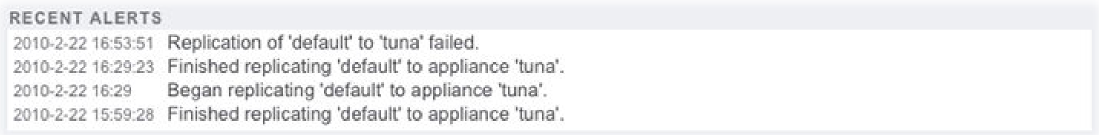

RECENT ALERTS

2010-2-22 16:53:51	Replication of 'default' to 'tuna' failed.
2010-2-22 16:29:23	Finished replicating 'default' to appliance 'tuna'.
2010-2-22 16:29	Began replicating 'default' to appliance 'tuna'.
2010-2-22 15:59:28	Finished replicating 'default' to appliance 'tuna'.

Summary of Pool Usage

The following table describes the pool usage properties.

Table 2-18 Summary of Pool Usage

Pool Usage	Description
Used	Space used by this pool including data, snapshots, and the first copy of deduplicated data, if applicable.
Available	Amount of remaining disk space available to the user. Refers to the amount of available space, excluding unused space that is reserved by projects and shares within a pool.

Table 2-18 (Cont.) Summary of Pool Usage

Pool Usage	Description
Compression	Current compression ratio achieved by this pool. If compression is disabled, the ratio is 1x.
Dedup	Current deduplicated data size in this pool, with the deduplication ratio in parenthesis. If deduplication is disabled, the size is 0 and the ratio is (1x).

Summary of Memory Usage

The following table describes the memory usage properties.

Table 2-19 Summary of Main Memory Usage

Main Memory (RAM) Usage	Description
Cache	Bytes in use by the filesystem cache to improve performance.
Unused	Bytes not currently in use. After booting, this value will decrease as space is used by the filesystem cache.
Mgmt	Bytes in use by the appliance management software.
Other	Bytes in use by miscellaneous operating system software.
Kernel	Bytes in use by the operating system kernel.

Note that users need the `analytics/component create+read` authorization to view the memory usage. Without this authorization, the memory details do not appear on the dashboard.

Disk Activity Dashboard

The activity area of the dashboard shows graphs of eight performance statistics by default. The example in this section shows disk operations/second. The statistical average is plotted in blue, and the maximum appears in light gray.

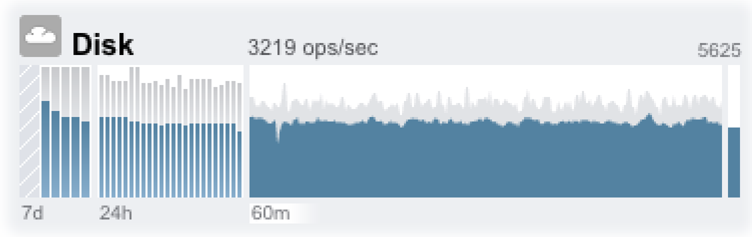

To go to the **Analytics** worksheet for an activity, click one of the four graphs (day, hour, minute, second) for the statistic you want to evaluate.

To view the average for each graph, mouse-over a graph and the average appears in the tooltip. The weather icon in the upper-left provides a report of activity according to thresholds you can customize for each statistic on the **Status Settings** screen.

Table 2-20 Summary of Statistic Graphs

Statistic Graph	Description
7-day graph (7d)	A bar chart, with each bar representing one day.
24-hour graph (24h)	A bar chart, with each bar representing one hour.
60-minute graph (60m)	A line plot, representing activity over one hour (also visible as the first one-hour bar in the 24-hour graph).
1-second graph	A line plot, representing instantaneous activity reporting.

The average for the selected plot is shown numerically above the graph. To change the average that appears, select the average you want, either 7d, 24h, or 60m.

The vertical scale of all graphs is printed on the top right, and all graphs are scaled to this same height. The height is calculated from the selected graph (plus a margin). The height will rescale based on activity in the selected graph, with the exception of utilization graphs which have a fixed height of 100 percent.

Since the height can rescale, 60 minutes of idle activity may look similar to 60 minutes of busy activity. Always check the height of the graphs before trying to interpret what they mean.

Understanding some statistics may not be obvious - you might wonder, for a particular appliance in your environment, whether 1000 NFSv3 ops/sec is considered busy or idle. This is where the 24-hour and 7-day plots can help, to provide historic data next to the current activity for comparison.

The plot height is calculated from the selected plot. By default, the 60-minute plot is selected. So, the height is the maximum activity during that 60-minute interval (plus a margin). To rescale all plots to span the highest activity during the previous 7 days, select 7d. This makes it easy to see how current activity compares to the last day or week.

The weather icon is intended to grab your attention when something is unusually busy or idle. To go to the weather threshold configuration page, click the weather icon. There is no good or bad threshold, rather the BUI provides a gradient of levels for each activity statistic. The statistics on which weather icons are based provide an *approximate* understanding for appliance performance that you should customize to your workload, as follows:

- Different environments have different acceptable levels for performance (latency), and so there is no one-size-fits-all threshold.
- The statistics on the Dashboard are based on operations/sec and bytes/sec, so you should use Analytics worksheets for an accurate understanding of system performance.

Dashboard CLI

A text version of the **Status: Dashboard** screen is available from the CLI by entering `status dashboard`:

```
hostname:> status dashboard
Data:
   pool_0:
      Used           497G bytes
      Avail          8.43T bytes
      State          online
      Compression    1x

Memory:
   Cache        30.1G bytes
   Unused       2.18G bytes
   Mgmt         343M bytes
   Other        474M bytes
   Kernel       38.9G bytes

Services:
   ad                disabled               smb              disabled
   dns               online                 ftp               disabled
   http              online                 identity          online
   idmap             online                 ipmp              online
   iscsi             online                 ldap              disabled
   ndmp              online                 nfs               online
   nis               online                 ntp               online
   routing           online                 scrk              maintenance
   snmp              online                 ssh               online
   tags              online                 vscan             online

Hardware:
   CPU               online                 Cards             online
   Disks             faulted                Fans              online
   Memory            online                 PSU               online

Activity:
   CPU              1 %util                 Sunny
   Disk            32 ops/sec               Sunny
   iSCSI            0 ops/sec               Sunny
   NDMP             0 bytes/sec             Sunny
   NFSv3            0 ops/sec               Sunny
   NFSv4            0 ops/sec               Sunny
   Network        13K bytes/sec             Sunny
   SMB              0 ops/sec               Sunny

Recent Alerts:
   2022-6-15 07:46: A cluster interconnect link has been restored.
```

The previous descriptions in this section apply, with the following differences:

- The activity plots are not rendered in text.
- The storage usage section will list details for all available pools in the CLI, whereas the BUI only has room to summarize one.

Separate views are available, for example `status activity show`:

```
hostname:> status activity show
Activity:
   CPU             10 %util                 Sunny
   Disk           478 ops/sec               Partly Cloudy
   iSCSI            0 ops/sec               Sunny
   NDMP             0 bytes/sec             Sunny
   NFSv3          681 ops/sec               Partly Cloudy
   NFSv4            0 ops/sec               Sunny
   Network      22.8M bytes/sec             Partly Cloudy
   SMB              0 ops/sec               Sunny
hostname:>
```

Running the Dashboard Continuously

You might experience browser memory issues if you leave the **Dashboard** screen open in a browser continuously (24x7). The browser will increase in size (memory leaks), and need to be closed and reopened. Browsers are fairly good at managing memory when browsing through different websites (and opening and closing tabs). The issue is that the Dashboard screen is left running and not closed, which opens and reopens images for the activity plots, thus degrading image rendering performance.

If you experience this problem while using the Firefox browser, disable the memory cache as follows:

1. Open **about:config**
2. Filter on **memory**.
3. Set **browser.cache.memory.enable = false**.

Status Dashboard Settings

The **Status: Settings** screen enables you to customize the **Status Dashboard**, including the statistics that appear and thresholds that indicate activity through the weather icons.

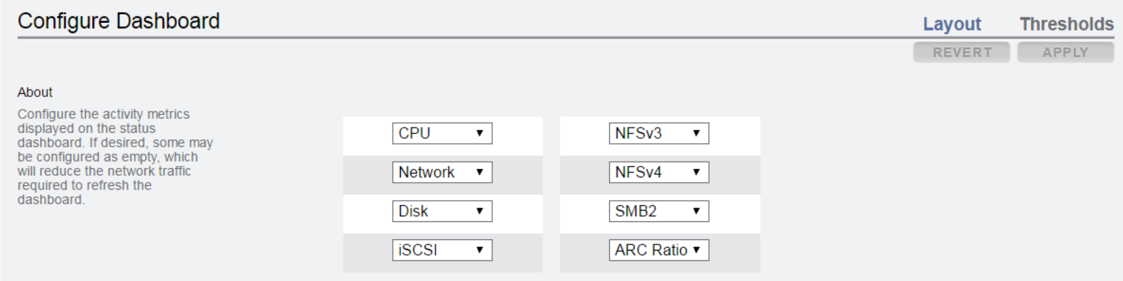

Configure Dashboard

About

Configure the activity metrics displayed on the status dashboard. If desired, some may be configured as empty, which will reduce the network traffic required to refresh the dashboard.

CPU ▾	NFSv3 ▾
Network ▾	NFSv4 ▾
Disk ▾	SMB2 ▾
iSCSI ▾	ARC Ratio ▾

Use the layout tab to select the graphs that appear in the dashboard activity area, as defined in the following table.

Table 2-21 Status Layout Settings

Name	Units	Description
<empty>	-	No graph will be displayed in this location.
CPU	utilization	Average cycles the appliance CPUs are busy. CPU cycles includes memory wait cycles.
ARC Ratio	utilization	Average ARC hit/miss percentage. A drop in the hit rate indicates a potential performance problem.
HTTP	operations/sec	Average number of HTTP operations.
Disk	operations/sec	Average number of operations to the physical storage devices.
iSCSI	operations/sec	Average number of iSCSI operations.
FC	operations/sec	Average number of Fibre Channel operations.
NDMP	bytes/sec	Average NDMP network bytes.
NFSv2	operations/sec	Average number of NFSv2 operations.
NFSv3	operations/sec	Average number of NFSv3 operations.
NFSv4.0	operations/sec	Average number of NFSv4.0 operations.
NFSv4.1	operations/sec	Average number of NFSv4.1 operations.
Network	bytes/sec	Average bytes/sec across all physical network interfaces.
SMB	operations/sec	Average number of SMB operations.

Table 2-21 (Cont.) Status Layout Settings

Name	Units	Description
SMB2	operations/sec	Average number of SMB2 operations.
SMB3	operations/sec	Average number of SMB3 operations.
FTP	bytes/sec	Average number of FTP bytes.
SFTP	bytes/sec	Average number of SFTP bytes.

Note that to reduce the network traffic required to refresh the **Dashboard**, configure some of the activity graphs as "<empty>".

Use the **Thresholds** screen to configure the dashboard activity weather icons. The defaults provided are based on heavy workloads and may not be suitable for your environment.

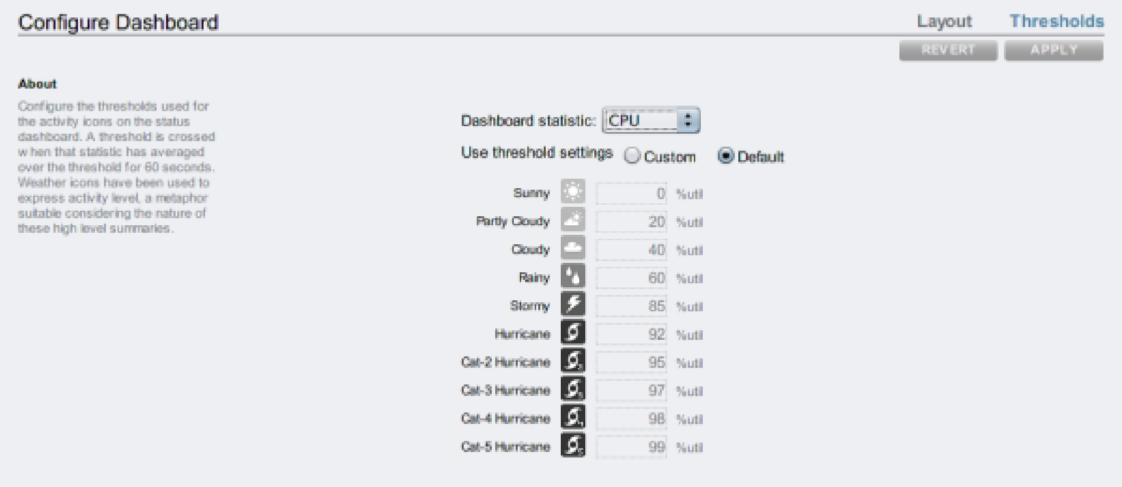

About

Configure the thresholds used for the activity icons on the status dashboard. A threshold is crossed w hen that statistic has averaged over the threshold for 60 seconds. Weather icons have been used to express activity level, a metaphor suitable considering the nature of these high level summaries.

Dashboard statistic: CPU

Use threshold settings ○ Custom ● Default

Sunny	0	%util
Partly Cloudy	20	%util
Cloudy	40	%util
Rainy	60	%util
Stormy	85	%util
Hurricane	92	%util
Cat-2 Hurricane	95	%util
Cat-3 Hurricane	97	%util
Cat-4 Hurricane	98	%util
Cat-5 Hurricane	99	%util

The weather icon that appears on the **Dashboard** is closest to the threshold value setting for the current activity, measured as a 60 second average. For example, if CPU utilization was at 41%, by default, the **Cloudy** weather icon would appear because its threshold is 40% (closest to the actual activity). Select the **Custom** radio button to configure thresholds and be sure to configure them in the order they appear on the screen.

The dashboard currently cannot be configured from the CLI. Settings saved in the BUI will apply to the dashboard that is visible from the CLI.

Changing the Displayed Activity Statistics

Use the following procedure to change the displayed activity statistics.

1. From the **Status** BUI menu, select **Settings**, then **Layout**.
2. From the drop-down menus, choose the statistics that you want to display on the **Dashboard**.
3. To save your choices, click **APPLY**.

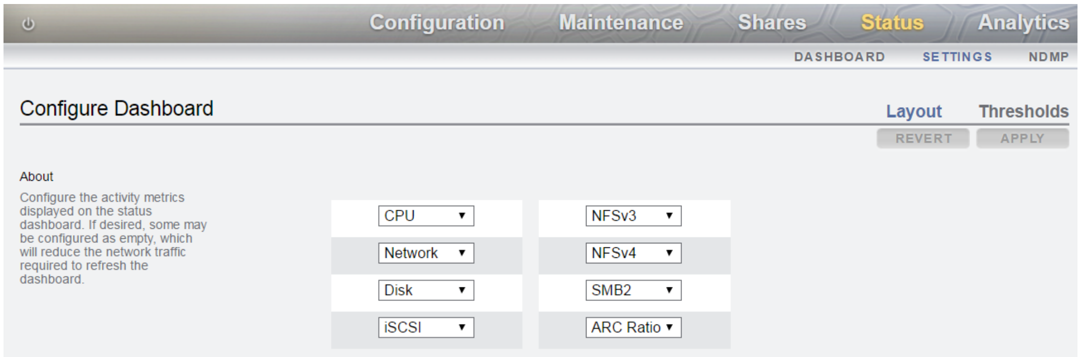

Configure Dashboard

About

Configure the activity metrics displayed on the status dashboard. If desired, some may be configured as empty, which will reduce the network traffic required to refresh the dashboard.

CPU	NFSv3
Network	NFSv4
Disk	SMB2
iSCSI	ARC Ratio

Changing the Activity Thresholds

Use the following procedure to change the activity thresholds.

1. From the **Status** BUI menu, select **Settings**, then **Thresholds**.
2. Choose the statistic to configure from the drop-down menu.
3. Click the **Custom** radio button.
4. Customize the values in the list, in the order that they appear. Some statistics will provide a **Units** drop-down menu, so that **Kilo/Mega/Giga** can be selected.
5. To save your configuration, click **APPLY**.

NDMP Status

When the NDMP service has been configured and is active, the **Status NDMP** page shows the NDMP devices and recent client activity. A green indicator shows that the device is online, and a gray indicator shows that the device is offline.

To resort the NDMB **Devices** list, click on the **Devices** column headings. To display details about a device, double-click on the device.

NDMP status is not available from the CLI.

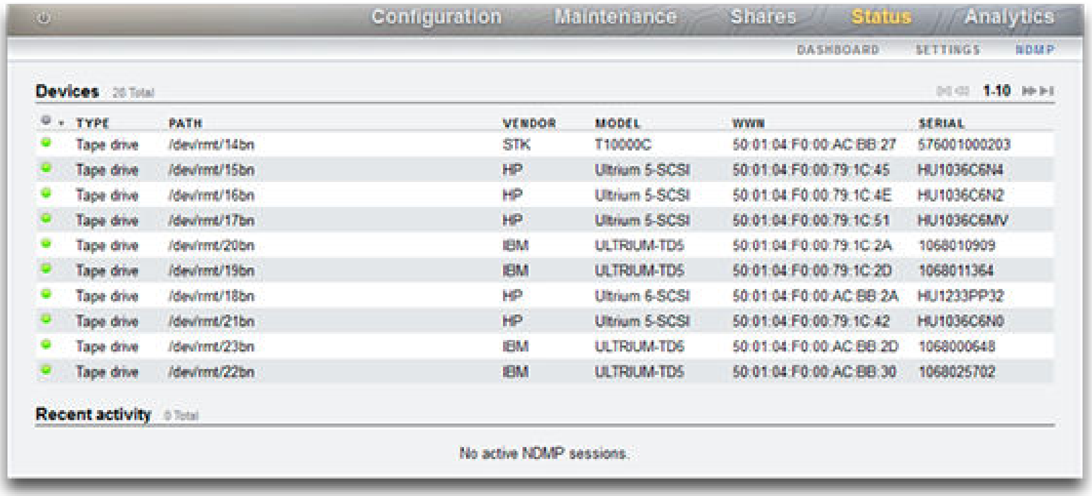

Devices 20 Total

	TYPE	PATH	VENDOR	MODEL	WWN	SERIAL
●	Tape drive	/dev/rmt/14bn	STK	T10000C	50:01:04:F0:00:AC:BB:27	576001000203
●	Tape drive	/dev/rmt/15bn	HP	Ultrium 5-SCSI	50:01:04:F0:00:79:1C:45	HU1036C6N4
●	Tape drive	/dev/rmt/16bn	HP	Ultrium 5-SCSI	50:01:04:F0:00:79:1C:4E	HU1036C6N2
●	Tape drive	/dev/rmt/17bn	HP	Ultrium 5-SCSI	50:01:04:F0:00:79:1C:51	HU1036C6MV
●	Tape drive	/dev/rmt/20bn	IBM	ULTRIUM-TD5	50:01:04:F0:00:79:1C:2A	1068010909
●	Tape drive	/dev/rmt/19bn	IBM	ULTRIUM-TD5	50:01:04:F0:00:79:1C:2D	1068011364
●	Tape drive	/dev/rmt/18bn	HP	Ultrium 6-SCSI	50:01:04:F0:00:AC:BB:2A	HU1233PP32
●	Tape drive	/dev/rmt/21bn	HP	Ultrium 5-SCSI	50:01:04:F0:00:79:1C:42	HU1036C6N0
●	Tape drive	/dev/rmt/23bn	IBM	ULTRIUM-TD6	50:01:04:F0:00:AC:BB:2D	1068000648
●	Tape drive	/dev/rmt/22bn	IBM	ULTRIUM-TD5	50:01:04:F0:00:AC:BB:30	1068025702

Recent activity 0 Total

No active NDMP sessions.

Table 2-22 NDMP Status - Devices

Field	Description	Examples
Type	Type of NDMP device	`Robot, Tape drive`
Path	Path of the NDMP device	`/dev/rmt/14bn`
Vendor	Device vendor name	`STK`
Model	Device model name	`T1000C`
WWN	World Wide Name	`50:01:04:F0:00:AC:BB:27`
Serial	Device serial number	`576001000203`

Table 2-23 NDMP Status - Recent Activity

Field	Description	Examples
ID	NDMP backup ID	`49`
Active	Backup currently active	`No`
Remote Client	NDMP client address and port	`192.168.1.219:4760`
Authenticated	Shows if the client has completed authentication yet	`Yes, No`
Data State	See Data State	`Active, Idle, ...`
Mover State	See Mover State	`Active, Idle, ...`
Current Operation	Current NDMP operation	`Backup, Restore, None`
Progress	A progress bar for this backup	

NDMP States

The NDMP **Data State** shows the state of the backup or restore operation. Possible values are:

- **Active** - The data is being backed up or restored.
- **Idle** - Backup or restore has not yet started or has already finished.
- **Connected** - Connection is established, but backup or restore has not yet begun.
- **Halted** - Backup or restore has finished successfully or has failed or aborted.
- **Listen** - Operation is waiting to receive a remote connection.

The NDMP **Mover State** shows the state of the NDMP device subsystem. Examples for tape devices are:

- **Active** - Data is being read from or written to the tape.
- **Idle** - Tape operation has not yet started or has already finished.
- **Paused** - Tape has reached the end or is waiting to be changed.
- **Halted** - Read/write operation has finished successfully or has failed or aborted.
- **Listen** - Operation is waiting to receive a remote connection.

Configuring Storage Area Network (SAN)

The **Storage Area Network (SAN)** configuration page lets you connect your Oracle ZFS Storage Appliance to your storage area network. A SAN comprises three basic components:

- A client that will access the storage on the network
- A storage appliance that will provide the storage on the network
- A network that will link the client to the storage

To configure SAN, use the following sections:

- Configuring FC Port Modes (BUI)
- Discovering FC Ports (BUI)
- Creating FC Initiator Groups (BUI)
- Associating a LUN with an FC Initiator Group (BUI)
- Changing FC Port Modes (CLI)
- Discovering FC Ports (CLI)
- Creating FC Initiator Groups (CLI)
- Associating a LUN with an FC Initiator Group (CLI)
- Scripting Aliases for Initiators and Initiator Groups (CLI)
- Configuring SAN iSCSI Initiators
- Creating an Analytics Worksheet (BUI)
- Adding an iSCSI Target with an Auto-generated IQN (CLI)
- Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)
- Adding an iSCSI Initiator with CHAP Authentication (CLI)
- Adding an iSCSI Target Group (CLI)
- Adding an iSCSI Initiator Group (CLI)
- Configuring SRP Target (BUI)
- Configuring SRP Targets (CLI)

To learn more about SAN, see the following sections:

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Configuring FC Port Modes (BUI)

Use the following procedure to configure FC port modes.

1. To use FC ports, set them to **Target** mode on the **Configuration: Storage Area Network (SAN)** screen of the BUI, using the drop-down menu shown in the following image. You must have root permissions to perform this action. Note that in a cluster configuration, you set ports to **Target** mode on each controller separately.

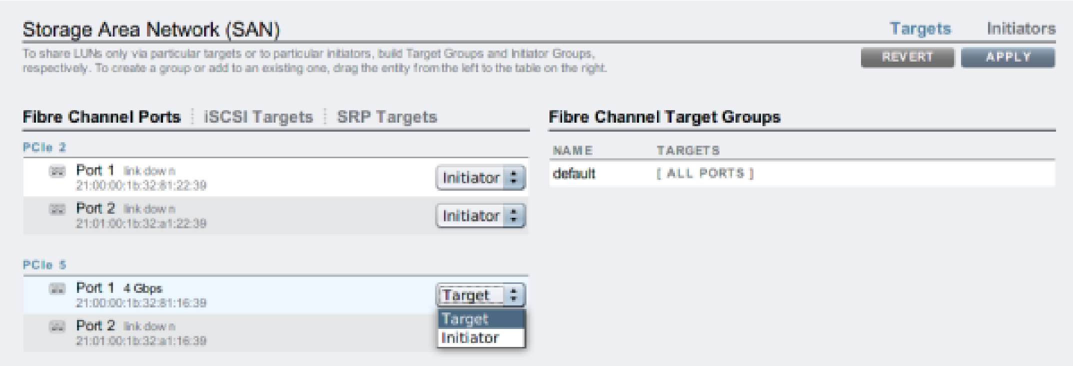

Storage Area Network (SAN)

To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right.

Fibre Channel Ports | iSCSI Targets | SRP Targets

PCIe 2

Port 1 link down 21:00:00:1b:32:81:22:39 — Initiator

Port 2 link down 21:01:00:1b:32:a1:22:39 — Initiator

PCIe 5

Port 1 4 Gbps 21:00:00:1b:32:81:16:39 — Target

Port 2 link down 21:01:00:1b:32:a1:16:39

Fibre Channel Target Groups

NAME	TARGETS
default	[ALL PORTS]

2. After setting desired ports to **Target**, click **APPLY**. A confirmation message will appear, notifying you that the appliance will reboot immediately. Confirm that you want to reboot.

3. When the appliance boots, the active FC targets appear with the FC active icon and, on mouse-over, the move icon ✥ appears.

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Discovering FC Ports (BUI)

Use the following procedure to discover FC ports.

1. Click the information icon ⓘ to view the **Discovered Ports** dialog box, where you can troubleshoot link errors.

2. In the **Discovered Ports** dialog box, click a **WWN** in the list to view associated link errors.

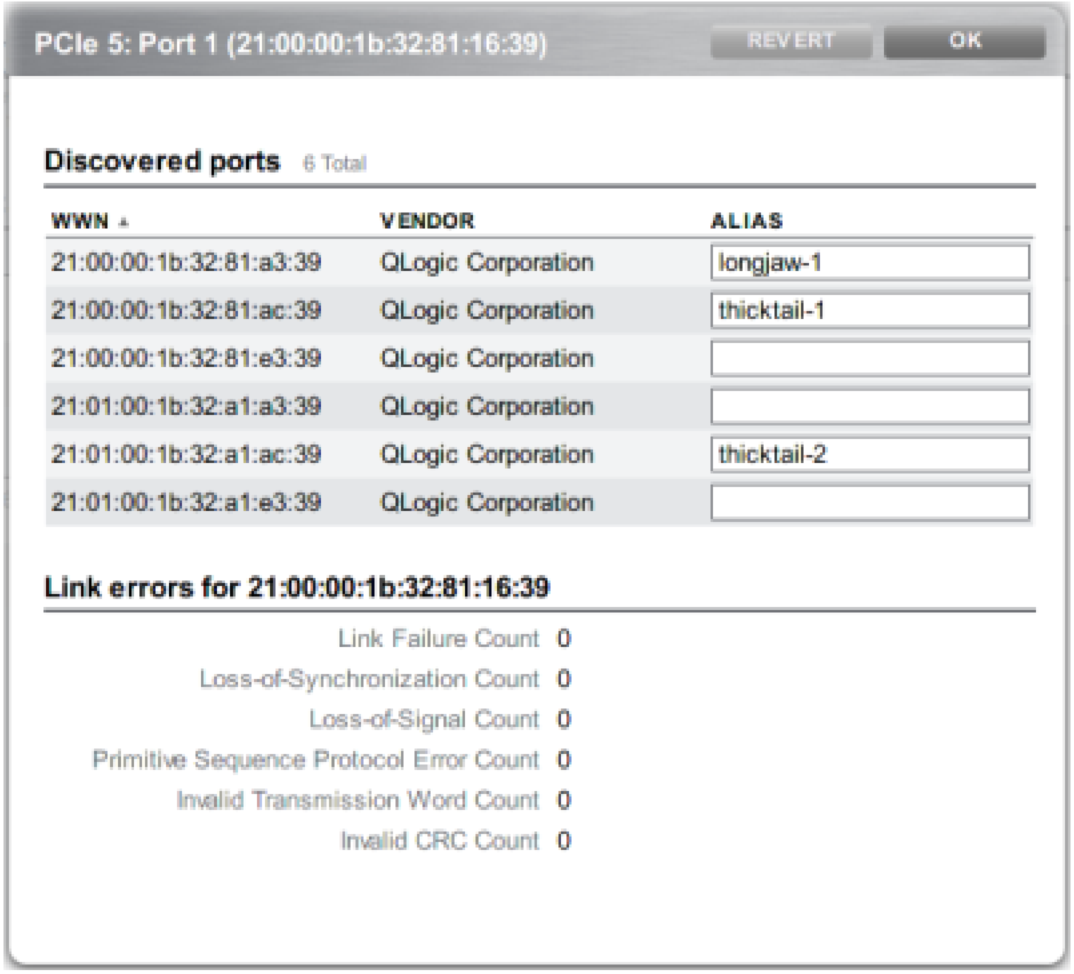

Discovered ports 6 Total

WWN ▴	VENDOR	ALIAS
21:00:00:1b:32:81:a3:39	QLogic Corporation	longjaw-1
21:00:00:1b:32:81:ac:39	QLogic Corporation	thicktail-1
21:00:00:1b:32:81:e3:39	QLogic Corporation	
21:01:00:1b:32:a1:a3:39	QLogic Corporation	
21:01:00:1b:32:a1:ac:39	QLogic Corporation	thicktail-2
21:01:00:1b:32:a1:e3:39	QLogic Corporation	

Link errors for 21:00:00:1b:32:81:16:39

Link Failure Count	0
Loss-of-Synchronization Count	0
Loss-of-Signal Count	0
Primitive Sequence Protocol Error Count	0
Invalid Transmission Word Count	0
Invalid CRC Count	0

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Creating FC Initiator Groups (BUI)

Use the following procedure to create FC initiator groups.

1. Create and manage initiator groups on the **Initiators** screen. Click the add icon ⊕ to view unaliased ports. Click a **WWN** in the list to add a meaningful alias in the **Alias** field.

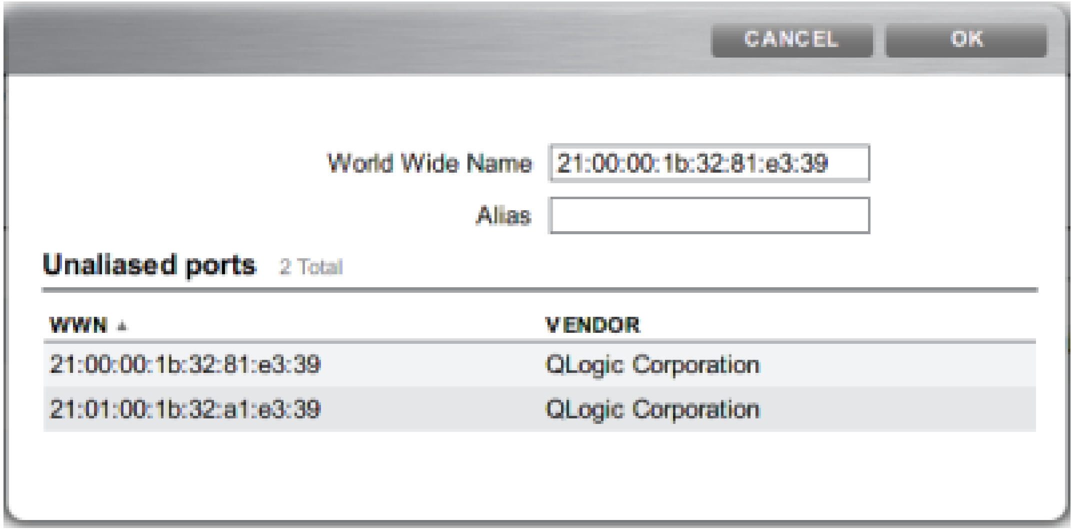

World Wide Name: 21:00:00:1b:32:81:e3:39

Alias:

Unaliased ports 2 Total

WWN	VENDOR
21:00:00:1b:32:81:e3:39	QLogic Corporation
21:01:00:1b:32:a1:e3:39	QLogic Corporation

2. On the **Initiators** page, drag initiators to the **FC Initiator Groups** list to create new groups or add to existing groups.

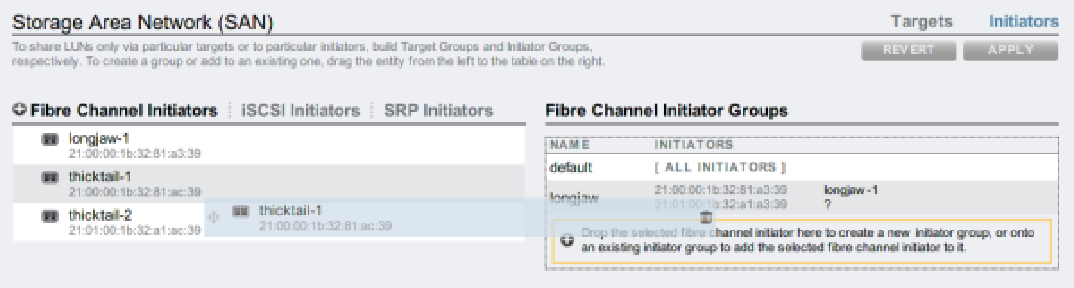

Storage Area Network (SAN)

To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right.

Fibre Channel Initiators:

- longjaw-1 21:00:00:1b:32:81:a3:39
- thicktail-1 21:00:00:1b:32:81:ac:39
- thicktail-2 21:01:00:1b:32:a1:ac:39
- thicktail-1 21:00:00:1b:32:81:ac:39

Fibre Channel Initiator Groups

NAME	INITIATORS	
default	[ALL INITIATORS]	
longjaw	21:00:00:1b:32:81:a3:39 21:01:00:1b:32:a1:a3:39	longjaw-1 ?

Drop the selected fibre channel initiator here to create a new initiator group, or onto an existing initiator group to add the selected fibre channel initiator to it.

3. Click **APPLY** to commit the new FC initiator group. Now you can create a LUN that has exclusive access to the client initiator group.

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Associating a LUN with an FC Initiator Group (BUI)

Use the following procedure to associate a LUN with an FC initiator group.

1. To create the LUN, roll-over the initiator group, and click the add LUN icon . The **Create LUN** dialog box appears with the associated initiator group selected.

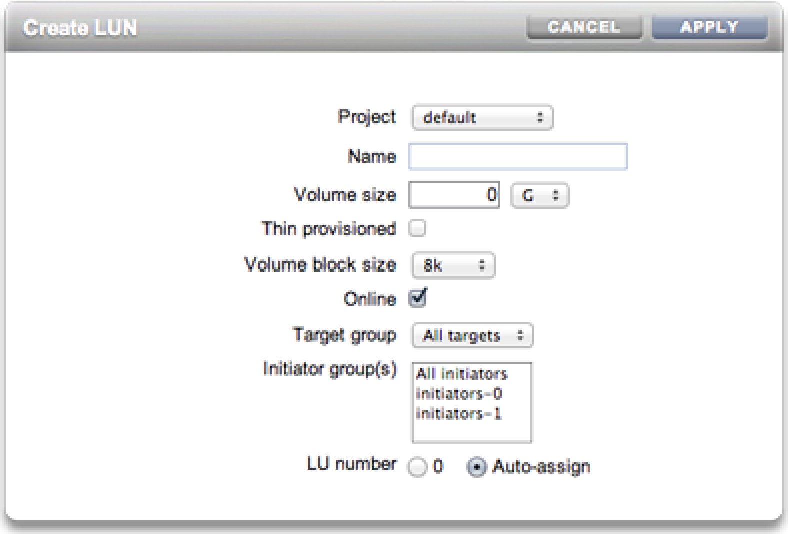

Project: default

Name:

Volume size: 0 G

Thin provisioned ☐

Volume block size: 8k

Online ☑

Target group: All targets

Initiator group(s): All initiators
initiators-0
initiators-1

LU number ○ 0 ● Auto-assign

2. Set the name and size, and click **APPLY** to add the LUN to the storage pool.

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Changing FC Port Modes (CLI)

To change FC port modes, use the following CLI commands:

```
hostname:configuration san fc targets> set targets="wwn.2101001B32A11639"
                      targets = wwn.2101001B32A11639 (uncommitted)
hostname:configuration san fc targets> commit
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Discovering FC Ports (CLI)

To discover FC ports, use the following CLI commands:

```
hostname:configuration san fc targets> show
Properties:
```

```
                      targets = wwn.2100001B32811639,wwn.2101001B32A12239
Targets:
NAME       MODE       WWN                     PORT                 SPEED
target-000 target     wwn.2100001B32811639    PCIe 5: Port 1       4 Gbit/s
target-001 initiator  wwn.2101001B32A11639    PCIe 5: Port 2       0 Gbit/s
target-002 initiator  wwn.2100001B32812239    PCIe 2: Port 1       0 Gbit/s
target-003 target     wwn.2101001B32A12239    PCIe 2: Port 2       0 Gbit/s
hostname:configuration san fc targets> select target-000
hostname:configuration san fc targets target-000> show
Properties:
                          wwn = wwn.2100001B32811639
                         port = PCIe 5: Port 1
                         mode = target
                        speed = 4 Gbit/s
             discovered_ports = 6
           link_failure_count = 0
           loss_of_sync_count = 0
         loss_of_signal_count = 0
         protocol_error_count = 0
        invalid_tx_word_count = 0
            invalid_crc_count = 0
Ports:
PORT      WWN                   ALIAS              MANUFACTURER
port-000  wwn.2100001B3281A339  longjaw-1          QLogic Corporation
port-001  wwn.2101001B32A1A339  longjaw-2          QLogic Corporation
port-002  wwn.2100001B3281AC39  thicktail-1        QLogic Corporation
port-003  wwn.2101001B32A1AC39  thicktail-2        QLogic Corporation
port-004  wwn.2100001B3281E339  <none>             QLogic Corporation
port-005  wwn.2101001B32A1E339  <none>             QLogic Corporation
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Creating FC Initiator Groups (CLI)

To create FC initiator groups, use the following CLI commands:

```
hostname:configuration san fc initiators> create
hostname:configuration san fc initiators (uncommitted)> set name=lefteye
hostname:configuration san fc initiators (uncommitted)> set
initiators=wwn.2101001B32A1AC39,wwn.2100001B3281AC39
hostname:configuration san fc initiators (uncommitted)> commit
hostname:configuration san fc initiators> list
GROUP     NAME
group-001 lefteye
         |
         +-> INITIATORS
             wwn.2101001B32A1AC39
             wwn.2100001B3281AC39
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Associating a LUN with an FC Initiator Group (CLI)

The following example demonstrates creating a LUN called `lefty` and associating it with the `fera` initiator group.

To associate a LUN with an FC initiator group, use the following CLI commands:

```
hostname:shares default> lun lefty
hostname:shares default/lefty (uncommitted)> set volsize=10
                      volsize = 10 (uncommitted)
hostname:shares default/lefty (uncommitted)> set initiatorgroup=fera
               initiatorgroup = default (uncommitted)
hostname:shares default/lefty (uncommitted)> commit
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Scripting Aliases for Initiators and Initiator Groups (CLI)

Refer to CLI Usage and Simple CLI Scripting and Batching Commands for information about how to modify and use the following example script.

To script aliases for initiators and initiator groups, use the following CLI commands:

```
script
     /*
      * This script creates both aliases for initiators and initiator
      * groups, as specified by the below data structure.  In this
      * particular example, there are five initiator groups, each of
      * which is associated with a single host (thicktail, longjaw, etc.),
      * and each initiator group consists of two initiators, each of which
      * is associated with one of the two ports on the FC HBA.  (Note that
      * there is nothing in the code that uses this data structure that
      * assumes the number of initiators per group.)
      */
      groups = {
               thicktail: {
                      'thicktail-1': 'wwn.2100001b3281ac39',
                      'thicktail-2': 'wwn.2101001b32a1ac39'
               },
               longjaw: {
```

```
                        'longjaw-1': 'wwn.2100001b3281a339',
                        'longjaw-2': 'wwn.2101001b32a1a339'
                },
                tecopa: {
                        'tecopa-1': 'wwn.2100001b3281e339',
                        'tecopa-2': 'wwn.2101001b32a1e339'
                },
                spinedace: {
                        'spinedace-1': 'wwn.2100001b3281df39',
                        'spinedace-2': 'wwn.2101001b32a1df39'
                },
                fera: {
                        'fera-1': 'wwn.2100001b32817939',
                        'fera-2': 'wwn.2101001b32a17939'
                }
        };
        for (group in groups) {
                initiators = [];
                for (initiator in groups[group]) {
                        printf('Adding %s for %s ... ',
                            groups[group][initiator], initiator);
                            try {
                                run('select alias=' + initiator);
                                printf('(already exists)\n');
                                run('cd ..');
                            } catch (err) {
                                if (err.code != EAKSH_ENTITY_BADSELECT)
                                        throw err;
                                run('create');
                                set('alias', initiator);
                                set('initiator', groups[group][initiator]);
                                run('commit');
                                printf('done\n');
                        }
                        run('select alias=' + initiator);
                        initiators.push(get('initiator'));
                        run('cd ..');
                }
                printf('Creating group for %s ... ', group);
                run('groups');
                try {
                        run('select name=' + group);
                        printf('(already exists)\n');
                        run('cd ..');
                } catch (err) {
                        if (err.code != EAKSH_ENTITY_BADSELECT)
                                throw err;
                        run('create');
                        set('name', group);
                        run('set initiators=' + initiators);
                        run('commit');
                        printf('done\n');
                }
                run('cd ..');
        }
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration

- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Creating an Analytics Worksheet (BUI)

To create an analytics worksheet for observing operations by initiator, complete the following steps.

1. Select the **Analytics** menu.
2. Click the add icon ⊕ for **Add Statistic**. A menu of all statistics appears.
3. Under the **Protocol** section of the menu, select **iSCSI operations**, then **Broken down by initiator**. A graph of the current operations by initiator appears.
4. To observe more detailed analytics, select the initiator from the field to the left of the graph, and click the drill-down icon . A menu of detailed analytics appears.

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Configuring SAN iSER Targets

In the BUI, iSER targets are managed as iSCSI targets on the **Configuration: Storage Area Network (SAN)** screen.

1. From the **Configuration** menu, select **Network**. To configure ibp(x) interfaces, select the ibp(x) interface (or ipmp) you want, and drag it to the **Datalinks** list to create the datalink.
2. Drag the datalink to the **Interfaces** list to create a new interface.

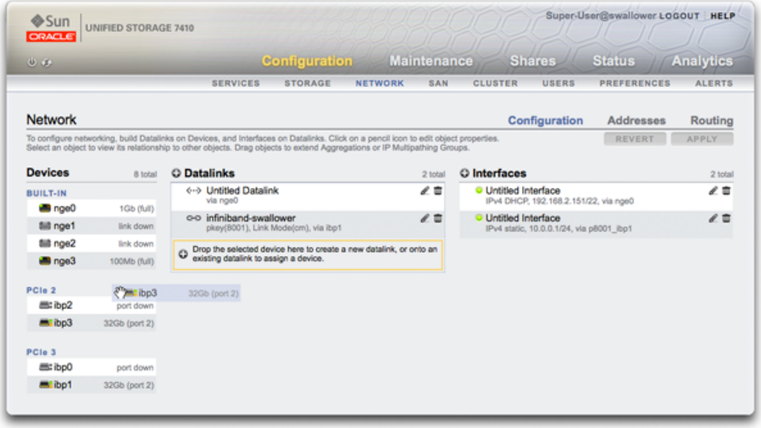

To configure networking, build Datalinks on Devices, and Interfaces on Datalinks. Click on a pencil icon to edit object properties. Select an object to view its relationship to other objects. Drag objects to extend Aggregations or IP Multipathing Groups.

Untitled Datalink
via nge0

infiniband-swallower
pkey(8001), Link Mode(cm), via ibp1

Drop the selected device here to create a new datalink, or onto an existing datalink to assign a device.

Untitled Interface
IPv4 DHCP, 192.168.2.151/22, via nge0

Untitled Interface
IPv4 static, 10.0.0.1/24, via p8001_ibp1

3. To create an iSER target, from the **Configuration** menu, select **SAN**, then **iSCSI Targets**.

4. To add a new iSER target with an alias, click the add icon ⊕ .
5. To create a target group, drag the target you just created to the **iSCSI Target Groups** list.

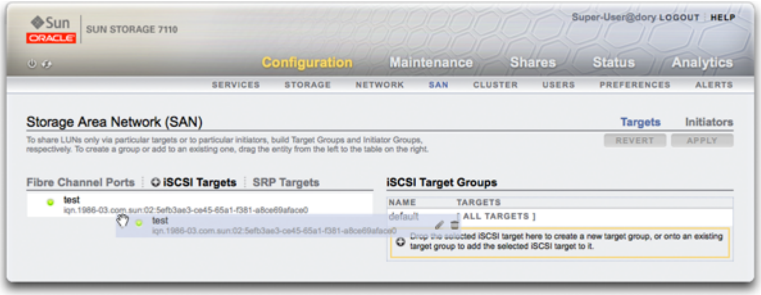

Storage Area Network (SAN)

To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right.

test
iqn.1986-03.com.sun:02:5efb3ae3-ce45-65a1-f381-a8ce69aface0

test
iqn.1986-03.com.sun:02:5efb3ae3-ce45-65a1-f381-a8ce69aface0

iSCSI Target Groups

NAME	TARGETS
default	[ALL TARGETS]

Drop the selected iSCSI target here to create a new target group, or onto an existing target group to add the selected iSCSI target to it.

6. To create an initiator, select **Initiators**, then **iSCSI Initiators**.
7. To add a new initiator, click the add icon ⊕ .
8. Enter the **Initiator IQN** and an alias, and click **OK**.

   Creating an initiator group is optional, but if you do not create a group, the LUN associated with the target will be available to all initiators.
9. To create a group, drag the initiator to the **iSCSI Initiator Groups** list.

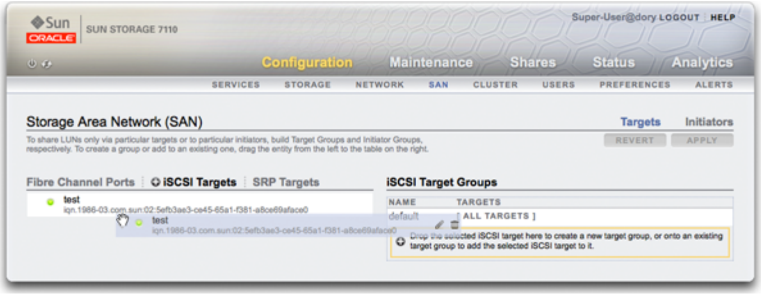

Storage Area Network (SAN)

To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right.

test
iqn.1986-03.com.sun:02:5efb3ae3-ce45-65a1-f381-a8ce69aface0

test
iqn.1986-03.com.sun:02:5efb3ae3-ce45-65a1-f381-a8ce69aface0

iSCSI Target Groups

NAME	TARGETS
default	[ALL TARGETS]

Drop the selected iSCSI target here to create a new target group, or onto an existing target group to add the selected iSCSI target to it.

10. To create a LUN, from the **Shares** menu, select **LUN**.
11. Click the add icon ⊕ and associate the new LUN with target or initiator groups you already created using the **Target group** and **Initiator group(s)** menus.

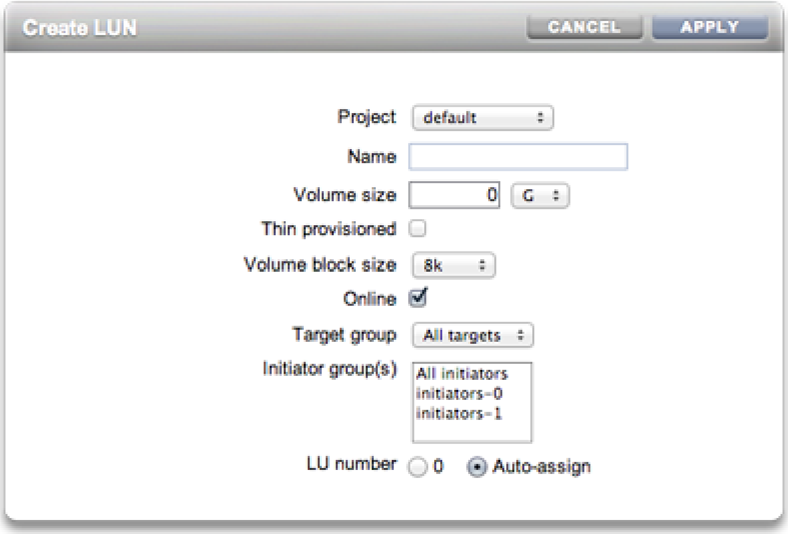

Project: default

Name:

Volume size: 0 G

Thin provisioned ☐

Volume block size: 8k

Online ☑

Target group: All targets

Initiator group(s): All initiators
initiators-0
initiators-1

LU number ○ 0 ● Auto-assign

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Adding an iSCSI Target with an Auto-generated IQN (CLI)

To add an iSCSI target with an auto-generated IQN, use the following CLI commands:

```
hostname:configuration san iscsi targets> create
hostname:configuration san iscsi targets target (uncommitted)> set alias="Target 0"
hostname:configuration san iscsi targets target (uncommitted)> set auth=none
hostname:configuration san iscsi targets target (uncommitted)> set interfaces=igb1
hostname:configuration san iscsi targets target (uncommitted)> commit
hostname:configuration san iscsi targets> list
TARGET     ALIAS
target-000 Target 0
           |
           +-> IQN
               iqn.1986-03.com.sun:02:daf0161f-9f5d-e01a-b5c5-e1efa9578416
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)

To add an iSCSI target with a specific IQN and RADIUS authentication, use the following CLI commands:

```
hostname:configuration san iscsi targets> create
hostname:configuration san iscsi targets target (uncommitted)> set alias="Target 1"
hostname:configuration san iscsi targets target (uncommitted)> set
iqn=iqn.2001-02.com.example:12345
hostname:configuration san iscsi targets target (uncommitted)> set auth=radius
hostname:configuration san iscsi targets target (uncommitted)> set interfaces=igb1
hostname:configuration san iscsi targets target (uncommitted)> commit
hostname:configuration san iscsi targets> list
TARGET     ALIAS
target-000 Target 0
           |
           +-> IQN
               iqn.1986-03.com.sun:02:daf0161f-9f5d-e01a-b5c5-e1efa9578416
target-001 Target 1
           |
           +-> IQN
               iqn.2001-02.com.acme:12345
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Adding an iSCSI Initiator with CHAP Authentication (CLI)

To add an iSCSI initiator with CHAP authentication, use the following CLI commands:

```
hostname:configuration san iscsi initiators> create
hostname:configuration san iscsi initiators initiator (uncommitted)> set
initiator=iqn.2001-02.com.example:initiator12345
hostname:configuration san iscsi initiators initiator (uncommitted)> set alias="Init 0"
hostname:configuration san iscsi initiators initiator (uncommitted)> set
chapuser=thisismychapuser
hostname:configuration san iscsi initiators initiator (uncommitted)> set
chapsecret=123456789012abc
hostname:configuration san iscsi initiators initiator (uncommitted)> commit
hostname:configuration san iscsi initiators> list
NAME          ALIAS
initiator-000 Init 0
              |
              +-> INITIATOR
                  iqn.2001-02.com.acme:initiator12345
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Adding an iSCSI Target Group (CLI)

To add an iSCSI target group, use the following CLI commands:

```
hostname:configuration san iscsi targets groups> create
hostname:configuration san iscsi targets group (uncommitted)> set name=tg0
hostname:configuration san iscsi targets group (uncommitted)> set
targets=iqn.2001-02.com.example:12345,iqn.1986-03.com.sample:02:daf0161f-9f5d-
e01a-b5c5-e1efa9578416
hostname:configuration san iscsi targets group (uncommitted)> commit
hostname:configuration san iscsi targets groups> list
GROUP     NAME
group-000 tg0
          |
          +-> TARGETS
              iqn.2001-02.com.acme:12345
              iqn.1986-03.com.sun:02:daf0161f-9f5d-e01a-b5c5-e1efa9578416
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Adding an iSCSI Initiator Group (CLI)

To add an iSCSI initiator group, use the following CLI commands:

```
hostname:configuration san iscsi initiators groups> create
hostname:configuration san iscsi initiators group (uncommitted)> set name=ig0
hostname:configuration san iscsi initiators group (uncommitted)> set
initiators=iqn.2001-02.com.example:initiator12345
hostname:configuration san iscsi initiators group (uncommitted)> commit
hostname:configuration san iscsi initiators groups> list
GROUP     NAME
group-000 ig0
          |
          +-> INITIATORS
             iqn.2001-02.com.acme:initiator12345
```

Related Topics

- Understanding SAN

- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Configuring SRP Targets (BUI)

This procedure describes the steps for configuring SRP targets.

1. Connect HCA ports to InfiniBand interfaces.
2. The targets are automatically discovered by the appliance.
3. To create the target group, from the **Configuration** menu, select **SAN**.
4. Click the **Target** link, and then click **SRP targets**.

   The **SRP targets** page appears.
5. To create the target group, use the move icon ✥ to drag a target to the **Target Groups** list.
6. Click **APPLY**.
7. Optional: To create an initiator and initiator group on the **Initiator** screen, click the add icon ⊕ and collect the GUID from the initiator, assign it a name, and drag it to the initiator group.
8. To create a LUN and associate it with the SRP target and initiators you created in the previous steps, select the **Shares** menu.
9. Click the **LUNs** link, and then click the LUN add icon ⊕ . Use the **Target Group** and **Initiator Group** menus in the **Create LUN** dialog box to select the SRP groups to associate with the LUN.

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Configuring SRP Targets (CLI)

To configure SRP targets, use the following CLI commands:

```
hostname:configuration san targets srp groups> create
hostname:configuration san targets srp group (uncommitted)> set name=targetSRPgroup
                         name = targetSRPgroup (uncommitted)
hostname:configuration san targets srp group (uncommitted)> set 
targets=eui.0002C903000489A4
                      targets = eui.0002C903000489A4 (uncommitted)
hostname:configuration san targets srp group (uncommitted)> commit
hostname:configuration san targets srp groups> list
GROUP     NAME
```

```
group-000 targetSRPgroup
         |
         +-> TARGETS
             eui.0002C903000489A4
```

The following example demonstrates how to create an SRP target group named `targetSRPgroup` using the CLI `configuration san targets srp groups` context:

Example 2-1 Creating a LUN associated with the Target SRP Group using the CLI

The following example shows how to create a LUN, and associate it with `targetSRPgroup` using the CLI `shares` context:

```
hostname:shares default> lun mylun
hostname:shares default/mylun (uncommitted)> set targetgroup=targetSRPgroup
                  targetgroup = targetSRPgroup (uncommitted)
hostname:shares default/mylun (uncommitted)> set volsize=10
                      volsize = 10 (uncommitted)
hostname:shares default/mylun (uncommitted)> commit
hostname:shares default> list
Filesystems:
NAME             SIZE    MOUNTPOINT
test             38K     /export/test
LUNs:
NAME              SIZE    GUID
mylun             10G     600144F0E9D19FFB00004B82DF490001
```

Related Topics

- Understanding SAN
- SAN Fibre Channel Configuration
- SAN iSCSI Configuration
- SAN iSER Target Configuration
- SAN SRP Configuration
- SAN Terminology

Understanding SAN

Fibre Channel, iSCSI, and SRP remain the same regardless of which protocol is used on the network. In some cases, the network may even be a cable between the initiator and the target, but in most cases, there is some type of switching involved.

Targets and initiators are configured by protocol. Refer to the documentation on a particular protocol for details (SAN Fibre Channel Configuration, iSCSI or SRP Configuration).

Target and initiator groups define sets of targets and initiators that can be associated with LUNs. A LUN that is associated with a target group can only be seen via the targets in the group. If a LUN is not explicitly associated with a target group, it is in the *default target group* and will be accessible via all targets, regardless of protocol. Similarly, a LUN can only be seen by the initiators in the group or groups to which it belongs. If a LUN is not explicitly associated with an initiator group, it is in the *default initiator group* and can be accessed by all initiators. While using the default initiator group can be useful for evaluation purposes, its use is discouraged since it may result in exposure of the LUN to unwanted or conflicting initiators.

To avoid possible LUN conflicts when an initiator belongs to multiple groups, configure initiators within all groups before associating groups with LUNs.

To configure targets, go to the **Configuration** BUI menu, select **SAN**, then **Fibre Channel** or **iSCSI** or **SRP** to navigate. Then configure the **Ports**, **Initiator**, and **Target Groups** controls.

To associate a LUN, from the **Shares** BUI menu, select **Shares**, then **Protocols**. Configure the **Target Group** and **Initiator Group** controls.

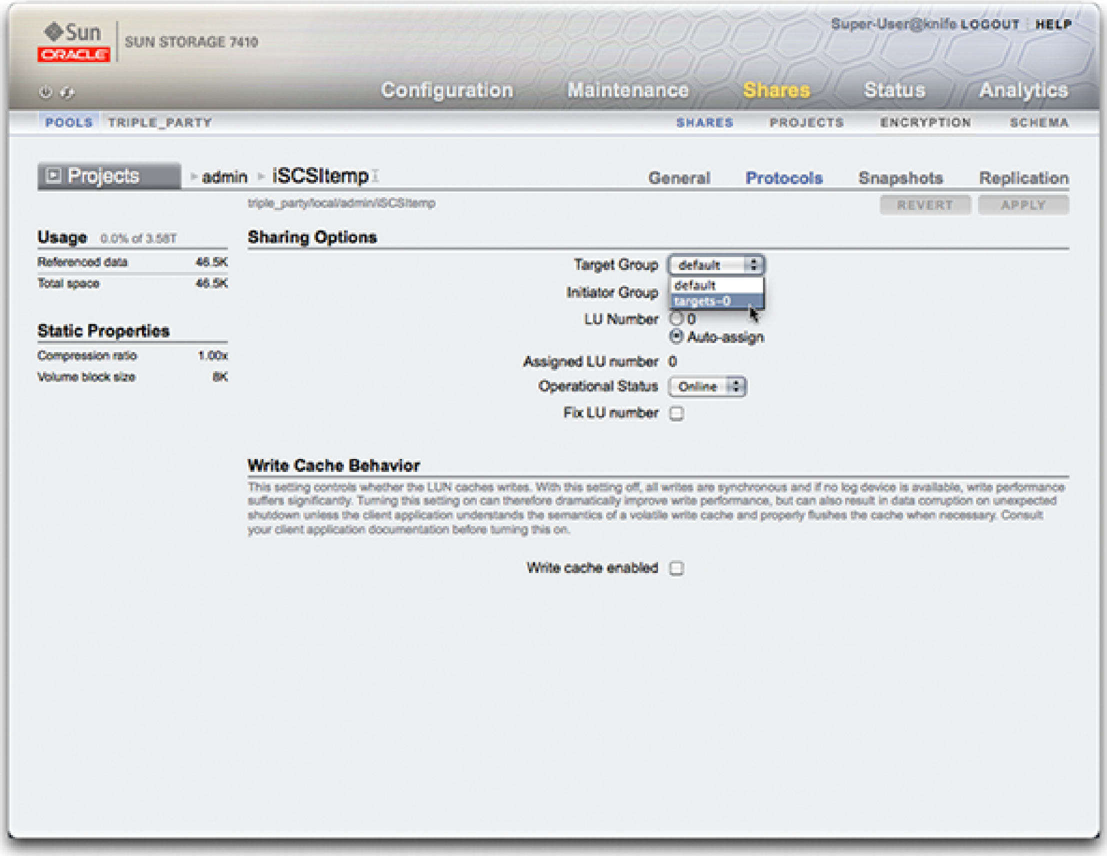

triple_party/local/admin/iSCSItemp

Usage 0.0% of 3.58T

Referenced data	46.5K
Total space	46.5K

Static Properties

Compression ratio	1.00x
Volume block size	8K

Sharing Options

Target Group: default
Initiator Group: default / targets-0
LU Number: ○ 0 ● Auto-assign
Assigned LU number: 0
Operational Status: Online
Fix LU number: □

Write Cache Behavior

This setting controls whether the LUN caches writes. With this setting off, all writes are synchronous and if no log device is available, write performance suffers significantly. Turning this setting on can therefore dramatically improve write performance, but can also result in data corruption on unexpected shutdown unless the client application understands the semantics of a volatile write cache and properly flushes the cache when necessary. Consult your client application documentation before turning this on.

Write cache enabled □

Use the `configuration san` context of the CLI to operate on targets and initiators by protocol type. Then, use the `shares` CLI context to create LUNs, and associate them with target and initiator groups.

Related Topics

- Configuring FC Port Modes (BUI)
- Discovering FC Ports (BUI)
- Creating FC Initiator Groups (BUI)
- Associating a LUN with an FC Initiator Group (BUI)
- Changing FC Port Modes (CLI)
- Discovering FC Ports (CLI)
- Creating FC Initiator Groups (CLI)
- Associating a LUN with an FC Initiator Group (CLI)
- Scripting Aliases for Initiators and Initiator Groups (CLI)
- Configuring SAN iSCSI Initiators

- Creating an Analytics Worksheet (BUI)
- Adding an iSCSI Target with an Auto-generated IQN (CLI)
- Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)
- Adding an iSCSI Initiator with CHAP Authentication (CLI)
- Adding an iSCSI Target Group (CLI)
- Adding an iSCSI Initiator Group (CLI)
- Configuring SRP Target (BUI)
- Configuring SRP Targets (CLI)

SAN Fibre Channel Configuration

Fibre Channel (FC) is a gigabit-speed networking technology used nearly exclusively as a transport for SCSI. FC is one of several block protocols supported by Oracle ZFS Storage Appliance; to share LUNs via FC, the appliance must be equipped with one or more optional FC cards.

By default, all FC ports are configured to be in target mode. If the appliance is used to connect to a tape SAN for backup, one or more ports must be configured in initiator mode. To configure a port for initiator mode, the appliance must be reset. Multiple ports can be configured for initiator mode simultaneously.

Each FC port is assigned a World Wide Name (WWN), and, as with other block protocols, FC targets may be grouped into SAN target and initiator groups, allowing port bandwidth to be dedicated to specific LUNs or groups of LUNs. After an FC port is configured as a target, the remotely discovered ports can be examined and verified.

Refer to the *Implementing Fibre Channel SAN Boot with Oracle ZFS Storage Appliance* technical brief at https://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/fc-sanboot-081412-pdf-1735984.pdf for details on FC SAN boot solutions using the appliance.

In a cluster, initiators will have two paths (or sets of paths) to each LUN: one path (or set of paths) will be to the head that has imported the storage associated with the LUN; the other path (or set of paths) will be to that head's clustered peer. The first path (or set of paths) is *active*; the second path (or set of paths) is *standby*. In the event of a takeover, the active paths will become unavailable, and the standby paths will (after a short time) be transitioned to be active, after which I/O will continue. This approach to multipathing is known as asymmetric logical unit access (ALUA) and, when coupled with an ALUA-aware initiator, allows cluster takeover to be transparent to higher-level applications.

Initiators are identified by their WWN. As with other block protocols, aliases can be created for initiators. To aid in creating aliases for FC initiators, a WWN can be selected from the WWNs of discovered ports. Also as with other block protocols, initiators can be collected into groups. When a LUN is associated with a specific initiator group, the LUN will only be visible to initiators in the group. In most FC SANs, LUNs will always be associated with the initiator group that corresponds to the system(s) for which the LUN has been created.

The appliance is an ALUA-compliant array. Properly configuring an FC initiator in an ALUA environment requires an ALUA-aware driver and may require initiator-specific tuning. See "Oracle ZFS Storage Appliance: How to set up Client Multipathing" (Doc ID 1628999.1) for more information.

FC performance can be observed via Analytics, whereby one can breakdown operations or throughput by initiator, target, or LUN:

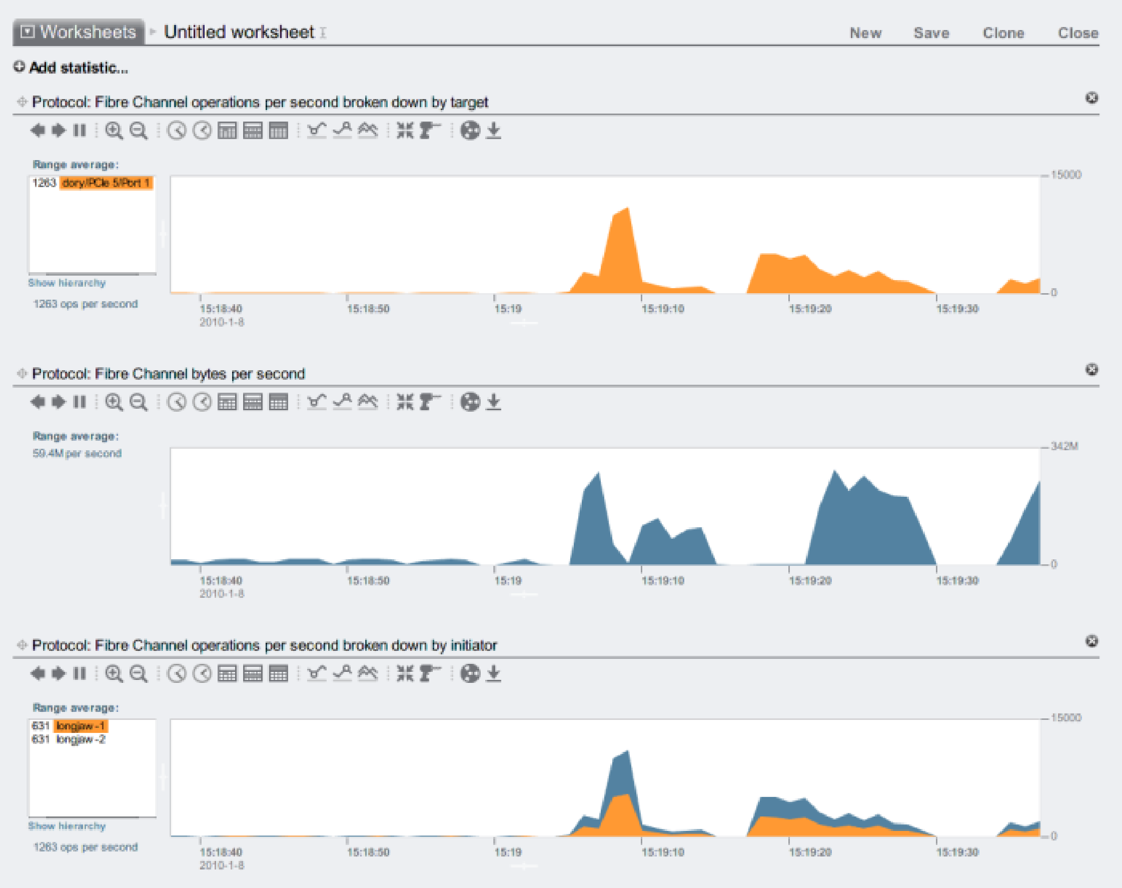

For operations, one can also breakdown by offset, latency, size and SCSI command, allowing one to understand not just the *what*, but the *how* and *why* of FC operations.

The appliance has been designed to utilize a global set of resources to service LUNs on each head. It is therefore not generally necessary to restrict queue depths on clients as the FC ports in the appliance can handle a large number of concurrent requests. Even so, there exists the remote possibility that these queues can be overrun, resulting in SCSI transport errors. Such queue overruns are often associated with one or more of the following:

- Overloaded ports on the front end - too many hosts associated with one FC port and/or too many LUNs accessed through one FC port
- Degraded appliance operating modes, such as a cluster takeover in what is designed to be an active-active cluster configuration

While the possibility of queue overruns is remote, it can be eliminated entirely if one is willing to limit queue depth on a per-client basis. To determine a suitable queue depth limit, one should take the number of target ports multiplied by the maximum concurrent commands per port (2048) and divide the product by the number of LUNs provisioned. To accommodate degraded operating modes, one should sum the number of LUNs across cluster peers to determine the number of LUNs, but take as the number of target ports the minimum of the two cluster peers. For example, in an active-active 7420 dual-headed cluster with one head having 2 FC ports and 100 LUNs and the other head having 4 FC ports and 28 LUNs, one should take the pessimal maximum queue depth to be two ports times 2048 commands divided by 100 LUNs plus 28 LUNs, or 32 commands per LUN.

Tuning the maximum queue depth is initiator specific, but on Oracle Solaris, this is achieved by adjusting the global variable `ssd_max_throttle`.

To troubleshoot link-level issues such as broken optics or a poorly seated cable, look at the error statistics for each FC port. If any number is either significantly non-zero or increasing, that may be an indicator that link-level issues have been encountered, and that link-level diagnostics should be performed.

Related Topics

- Configuring FC Port Modes (BUI)
- Discovering FC Ports (BUI)
- Creating FC Initiator Groups (BUI)
- Associating a LUN with an FC Initiator Group (BUI)
- Changing FC Port Modes (CLI)
- Discovering FC Ports (CLI)
- Creating FC Initiator Groups (CLI)
- Associating a LUN with an FC Initiator Group (CLI)
- Scripting Aliases for Initiators and Initiator Groups (CLI)
- Configuring SAN iSCSI Initiators
- Creating an Analytics Worksheet (BUI)
- Adding an iSCSI Target with an Auto-generated IQN (CLI)
- Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)
- Adding an iSCSI Initiator with CHAP Authentication (CLI)
- Adding an iSCSI Target Group (CLI)
- Adding an iSCSI Initiator Group (CLI)
- Configuring SRP Target (BUI)
- Configuring SRP Targets (CLI)

SAN iSCSI Configuration

Internet SCSI is one of several block protocols supported by Oracle ZFS Storage Appliance for sharing SCSI-based storage. When using the iSCSI protocol, the target portal refers to the unique combination of an IP address and TCP port number by which an initiator can contact a target.

When using the iSCSI protocol, a target portal group is a collection of target portals. Target portal groups are managed transparently; each network interface has a corresponding target portal group with that interface's active addresses. Binding a target to an interface advertises that iSCSI target using the portal group associated with that interface.

> **Note:**
>
> Multiple connections per session are not supported.

An IQN (iSCSI qualified name) is the unique identifier of a device in an iSCSI network. iSCSI uses the form `iqn.date.authority:uniqueid` for IQNs. For example, the appliance may use the IQN: `iqn.1986-03.com.example:02:c7824a5b-f3ea-6038-c79d-ca443337d92c` to identify one of its iSCSI targets. This name shows that this is an iSCSI device built by a company registered in March of 1986. The naming authority is just the DNS name of the company reversed, in this case, `com.example`. Everything following is a unique ID that Oracle uses to identify the target.

Table 2-24 iSCSI Target Properties

Target Property	Description
Target IQN	The IQN for this target; the IQN can be manually specified or auto-generated
Alias	A human-readable nickname for this target
Authentication mode	One of None, CHAP, or RADIUS
CHAP name	If CHAP authentication is used, the CHAP username
CHAP secret	If CHAP authentication is used, the CHAP secret
Network interfaces	The interfaces whose target portals are used to export this target

In addition to those properties, the BUI indicates whether a target is online or offline.

Table 2-25 Target Status Icons

Icon	Description
▟	Target is online
●	Target is offline

On clustered platforms, targets which have at least one active interface on that cluster node will be online. Take care when assigning interfaces to targets; a target may be configured to use portal groups on disjoint head nodes. In that situation, the target will be online on both heads yet will export different LUNs depending on the storage owned by each head node. As network interfaces migrate between cluster heads as part of takeover/failback or ownership changes, iSCSI targets will move online and offline as their respective network interfaces are imported and exported.

Targets which are bound to an IPMP interface will be advertised only via the addresses of that IPMP group. That target will not be reachable via that group's test addresses. Targets bound to interfaces built on top of a LACP aggregation will use the address of that aggregation. If a LACP aggregation is added to an IPMP group, a target can no longer use that aggregation's interface, as that address will become an IPMP test address.

Related Topics

- Configuring FC Port Modes (BUI)
- Discovering FC Ports (BUI)
- Creating FC Initiator Groups (BUI)
- Associating a LUN with an FC Initiator Group (BUI)
- Changing FC Port Modes (CLI)
- Discovering FC Ports (CLI)
- Creating FC Initiator Groups (CLI)
- Associating a LUN with an FC Initiator Group (CLI)

- Scripting Aliases for Initiators and Initiator Groups (CLI)
- Configuring SAN iSCSI Initiators
- Creating an Analytics Worksheet (BUI)
- Adding an iSCSI Target with an Auto-generated IQN (CLI)
- Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)
- Adding an iSCSI Initiator with CHAP Authentication (CLI)
- Adding an iSCSI Target Group (CLI)
- Adding an iSCSI Initiator Group (CLI)
- Configuring SRP Target (BUI)
- Configuring SRP Targets (CLI)

SAN iSCSI Initiator Configuration

iSCSI initiators have the following configurable properties.

Table 2-26 SAN iSCSI Initiator Properties

Property	Description
Initiator IQN	The IQN for this initiator
Alias	A human-readable nickname for this initiator
Use CHAP	Enables or disables CHAP authentication
CHAP name	If CHAP authentication is used, the CHAP username
CHAP secret	If CHAP authentication is used, the CHAP secret

When planning your iSCSI client configuration, you will need the following information:

- Which initiators (and their IQNs) will be accessing the SAN?
- If you plan on using CHAP authentication, what CHAP credentials does each initiator use?
- How many iSCSI disks (LUNs) are required, and how big should they be?
- Do the LUNs need to be shared between multiple initiators?

To allow the appliance to perform CHAP authentication using RADIUS, the following pieces of information must match:

- The appliance must specify the address of the RADIUS server and a secret to use when communicating with this RADIUS server
- The RADIUS server (for example, in its clients file) must have an entry giving the address of this appliance and specifying the same secret as earlier
- The RADIUS server (for example, in its users file) must have an entry giving the CHAP name and matching CHAP secret of each initiator
- If the initiator uses its IQN name as its CHAP name (the recommended configuration), then the appliance does not need a separate initiator entry for each initiator box; the RADIUS server can perform all authentication steps.
- If the initiator uses a separate CHAP name, then the appliance must have an initiator entry for that initiator that specifies the mapping from IQN name to CHAP name. This initiator entry does NOT need to specify the CHAP secret for the initiator.

For tips on troubleshooting common iSCSI misconfiguration, see iSCSI.

iSCSI performance can be observed via Analytics, whereby one can break down operations or throughput by initiator, target, or LUN.

Related Topics

- Configuring FC Port Modes (BUI)
- Discovering FC Ports (BUI)
- Creating FC Initiator Groups (BUI)
- Associating a LUN with an FC Initiator Group (BUI)
- Changing FC Port Modes (CLI)
- Discovering FC Ports (CLI)
- Creating FC Initiator Groups (CLI)
- Associating a LUN with an FC Initiator Group (CLI)
- Scripting Aliases for Initiators and Initiator Groups (CLI)
- Configuring SAN iSCSI Initiators
- Creating an Analytics Worksheet (BUI)
- Adding an iSCSI Target with an Auto-generated IQN (CLI)
- Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)
- Adding an iSCSI Initiator with CHAP Authentication (CLI)
- Adding an iSCSI Target Group (CLI)
- Adding an iSCSI Initiator Group (CLI)
- Configuring SRP Target (BUI)
- Configuring SRP Targets (CLI)

SAN SRP Configuration

SCSI RDMA Protocol (SRP) is a protocol supported by Oracle ZFS Storage Appliance for sharing SCSI-based storage over a network that provides RDMA services (that is, InfiniBand).

SRP ports are shared with other IB port services, such as IPoIB and RDMA. The SRP service may only operate in target mode. SRP targets have the following configurable properties.

Table 2-27 SRP Target Properties

Property	Description
Target EUI	The Extended Unique Identifier (EUI) for this target. The EUI is automatically assigned by the system and is equal to the HCA GUID over which the SRP port service is running.
Alias	A human-readable nickname for this target.

In addition to those properties, the BUI indicates whether a target is online or offline:

Table 2-28 SRP Target Status Icons

Icon	Description
▟	Target is online
●	Target is offline

On clustered platforms, peer targets should be configured into the same target group for highly available (multipathed) configurations. SRP multipathed I/O is an initiator-side configuration option.

SRP initiators have the following configurable properties.

Table 2-29 SRP Initiator Properties

Property	Description
Initiator EUI	The EUI for this initiator
Alias	A human-readable nickname for this initiator

SRP performance can be observed via Analytics, whereby one can break down operations or throughput by initiator or target.

Related Topics

- Configuring FC Port Modes (BUI)
- Discovering FC Ports (BUI)
- Creating FC Initiator Groups (BUI)
- Associating a LUN with an FC Initiator Group (BUI)
- Changing FC Port Modes (CLI)
- Discovering FC Ports (CLI)
- Creating FC Initiator Groups (CLI)
- Associating a LUN with an FC Initiator Group (CLI)
- Scripting Aliases for Initiators and Initiator Groups (CLI)
- Configuring SAN iSCSI Initiators
- Creating an Analytics Worksheet (BUI)
- Adding an iSCSI Target with an Auto-generated IQN (CLI)
- Adding an iSCSI Target with a Specific IQN and RADIUS Authentication (CLI)
- Adding an iSCSI Initiator with CHAP Authentication (CLI)
- Adding an iSCSI Target Group (CLI)
- Adding an iSCSI Initiator Group (CLI)
- Configuring SRP Target (BUI)
- Configuring SRP Targets (CLI)

SAN Terminology

To configure Oracle ZFS Storage Appliance to operate on a SAN, you should understand some basic SAN terms:

Table 2-30 SAN Terminology

Term	Description
SCSI Target	A *SCSI Target* is a storage system end-point that provides a service of processing SCSI commands and I/O requests from an initiator. A SCSI Target is created by the storage system's administrator, and is identified by unique addressing methods. A SCSI Target, once configured, consists of zero or more logical units.
SCSI Initiator	A *SCSI Initiator* is an application or production system end-point that is capable of initiating a SCSI session, sending SCSI commands and I/O requests. SCSI Initiators are also identified by unique addressing methods (See SCSI Target).
Logical Unit	A *Logical Unit* is a term used to describe a component in a storage system. Uniquely numbered, this creates what is referred to as a Logical Unit Number or LUN. A storage system, being highly configurable, may contain many LUNS. These LUNs, when associated with one or more SCSI Targets, forms a unique SCSI device, a device that can be accessed by one or more SCSI Initiators.
iSCSI	*Internet SCSI (iSCSI)* is a protocol for sharing SCSI-based storage over IP networks. The appliance supports the SCSI-3 Persistent Reservations specification.
iSER	*iSCSI Extension for RDMA (iSER)* is a protocol that maps the iSCSI protocol over a network that provides RDMA services (that is, InfiniBand). The iSER protocol is transparently selected by the iSCSI subsystem, based on the presence of correctly configured IB hardware. In the CLI and BUI, all iSER-capable components (targets and initiators) are managed as iSCSI components.
FC	*Fibre Channel (FC)* is a protocol for sharing SCSI based storage over a storage area network (SAN), consisting of fiber-optic cables, FC switches and HBAs. The appliance supports 4GB and 8GB Fibre Channel Arbitrated Loop (FC-AL) topologies.
SRP	*SCSI RDMA Protocol (SRP)* is a protocol for sharing SCSI-based storage over a network that provides RDMA services (that is, InfiniBand).
IQN	An *iSCSI qualified name (IQN)* is the unique identifier of a device in an iSCSI network. iSCSI uses the form `iqn.date.authority:uniqueid` for IQNs. For example, the appliance may use the IQN: `iqn.1986-03.com.example:02:c7824a5b-f3ea-6038-c79d-ca443337d92c` to identify one of its iSCSI targets. This name shows that this is an iSCSI device built by a company registered in March of 1986. The naming authority is just the DNS name of the company reversed, in this case, `com.example`. Everything following is a unique ID that the company uses to identify the target.
Target Portal	When using the iSCSI protocol, the *Target Portal* refers to the unique combination of an IP address and TCP port number by which an initiator can contact a target.
Target Portal Group	When using the iSCSI protocol, a *Target Portal Group* is a collection of target portals. Target portal groups are managed transparently; each network interface has a corresponding target portal group with that interface's active addresses. Binding a target to an interface advertises that iSCSI target using the portal group associated with that interface.
CHAP	*Challenge-handshake authentication protocol (CHAP)* is a security protocol that can authenticate a target to an initiator, an initiator to a target, or both.

Table 2-30 (Cont.) SAN Terminology

Term	Description
RADIUS	*RADIUS* is a system for using a centralized server to perform authentication on behalf of storage nodes.
Target Group	A set of targets. LUNs are exported over all the targets in one specific *Target Group.*
Initiator Group	A set of initiators. When an *Initiator Group* is associated with a LUN, only initiators from that group may access the LUN.
Target	A storage system end-point that provides a service of processing SCSI commands and I/O requests from an initiator. A *Target* is created by the storage system administrator, and is identified by unique addressing methods. A target, once configured, consists of zero or more logical units.
Initiator	An application or production system end-point that is capable of initiating a SCSI session, sending SCSI commands and I/O requests. *Initiators* are also identified by unique addressing methods.

Each LUN has several properties which control how the volume is exported. For more information, see Protocols.

Configuring Users

This section describes how to create Oracle ZFS Storage Appliance users, including how to grant authorizations to users, and how to use roles to manage authorizations.

To configure users and roles, use the following sections:

- Adding an Administrator or User - BUI, CLI
- Changing a User Password - BUI, CLI
- Editing Exceptions for a User - BUI, CLI
- Adding a Role - BUI, CLI
- Editing Authorizations for a Role - BUI, CLI
- Adding a User Who Can Only View the Dashboard - BUI, CLI
- Determining the Current Logged-in User

To understand users and roles, see the following sections:

- Understanding Users and Roles
- User Authorizations
- Managing User Properties

Adding an Administrator or User (BUI)

Use the following procedure to create a user with or without the administrator role.

1. From the **Configuration** menu, select **Users**.
2. Click the add icon ⊕ next to **Users**.

   See also "Alternative Method" following this procedure.

3. In the **Add User** dialog box, choose the appropriate type of user from the **Type** drop-down menu.

   For descriptions of user types, see Understanding Users and Roles.

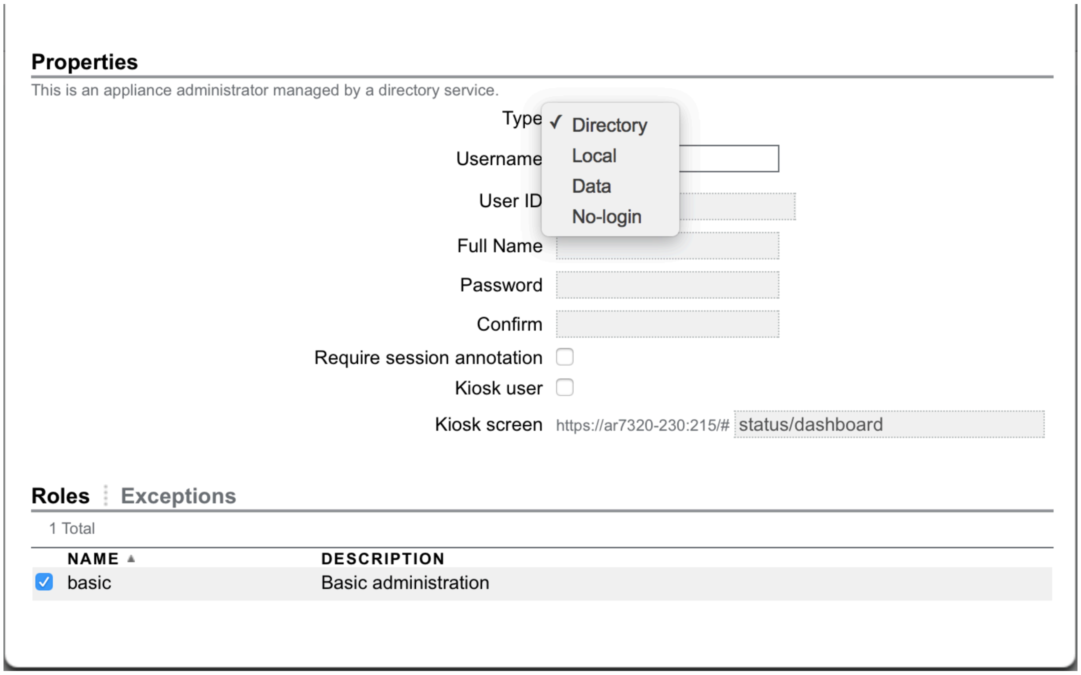

Properties

This is an appliance administrator managed by a directory service.

Type ✓ Directory

Username

User ID

Full Name

Password

Confirm

Require session annotation □

Kiosk user □

Kiosk screen https://ar7320-230:215/# status/dashboard

Roles | Exceptions

1 Total

	NAME ▲	DESCRIPTION
☑	basic	Basic administration

4. Enter values for properties.

   A **Username** value is required. For help with setting the **Username** value, see Understanding Users and Roles.

   - If you selected **Directory** for the user type, the **User ID** and **Password** are managed automatically by the directory service.
   - If you selected **Local**, **Data**, or **No-login** for the user type, then a **User ID** is required. If you do not set the **User ID**, a user ID is automatically assigned. For help with setting the **User ID** value, see "User Properties" in Managing User Properties.
   - If you selected **Local** or **Data** for the user type, then a **Password** is required.

   If you do not provide a value for **Full Name**, then the BUI will not show the identity of the current user. See Determining the Current Logged-in User.

5. Optional: For **Local** and **Directory** users, assign roles.

   Click the **Roles** tab. **Local** and **Directory** users have the `basic` role by default.

   Roles that are listed in the **Roles** section of **Configuration: Users** are available to choose. Click the check boxes for the roles that you want this user to have.

6. Optional: For **Local** and **Directory** users, add authorizations.

   See "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

   Click the **Exceptions** tab. Iterate the following steps until you have added all of the authorizations that you want this user to have:

   a. Select a **Scope**.

      Any filters that are available for this scope appear below the **Scope** selector.

   b. Specify filters for the scope as necessary.
   c. Click the check box for each authorization to add.
   d. Click **ADD** in the **Exceptions** section.

      The authorizations are listed at the bottom of the **Exceptions** section.

   Note that these authorizations can also be used to exclude authorizations that are granted to this user in a role. If you assign authorizations that have a more limited (more narrowly filtered) scope than the same authorizations that are granted in a role, then this user will only have the authorizations for the more limited scope.

7. Click **ADD** at the top of the dialog box.

   The new user appears in the **Users** list.

Alternative Method

To create a new user of the same type as an existing user and with the same roles and authorizations assigned, hover over the entry for the existing user and click the clone icon ⊞. Provide a **Username** and **Full Name**. If the type of the user that you are cloning is **Local** or **Data**, set a **Password**. Click **ADD** at the top of the **Clone User** dialog box.

Related Topics

- Understanding Users and Roles
- Managing User Properties
- User Authorizations
- Setting Appliance Preferences

Adding an Administrator or User (CLI)

Use the following procedure to create a user with or without the administrator role.

1. Go to `configuration users`.

   ```
   hostname:> configuration users
   ```

2. Enter a user type followed by a username.

   ```
   hostname:configuration users> type username
   ```

   For descriptions of user types and for help with setting the username value, see Understanding Users and Roles.

   If you specify `directory` for the user type, the user is immediately configured because no additional information is needed.

   ```
   hostname:configuration users> directory NISorLDAPorAD-username
   hostname:configuration users>
   ```

   If you specify `local`, `data`, or `nologin` for the user type, you are prompted to set properties.

   ```
   hostname:configuration users> local username
   hostname:configuration users username (uncommitted)>
   ```

3. Set properties for `local`, `data`, and `nologin` users.

   a. Enter `get` to list the properties to set.

```
hostname:configuration users username (uncommitted)> get
Properties:
                       logname = username
                          type = local
                           uid = (unset)
                      fullname = (unset)
              initial_password = (unset)
            require_annotation = false
```

   b. Set required properties.

   - For `local`, `data`, and `nologin` types, a user ID is required. If you do not set `uid` explicitly, a `uid` is automatically assigned. For help with setting the `uid` value, see "User Properties" in Managing User Properties.
   - For `local` and `data` types, you must set `initial_password`.

   c. Enter `commit`.

4. Enter `show`.

   The new user is listed.

5. Optional: Set additional properties and preferences.

   a. Select the new user.

   b. Enter `show` to see what you can set.

   You might see additional properties that you can set. For descriptions of properties, see "User Properties" in Managing User Properties.

   c. Enter preferences, and then enter `show` to see what preferences you can set. See Setting Appliance Preferences.

6. Optional: For local and directory users, assign additional roles.

   Users of type `local` or `directory` have the `basic` role by default.

   Roles that are listed in `configuration roles` are available to choose.

   a. Enter the following command to add a role for this user:

```
hostname:configuration users username> set roles=basic,additional_role
                         roles = basic,additional_role (uncommitted)
```

   b. Enter `commit`.

7. Optional: For local and directory users, assign additional authorizations.

   See the table for "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

   a. Select the new user.

   b. Enter `exceptions`.

```
hostname:configuration users username> exceptions
```

   c. Iterate the following steps until you have added all of the authorizations that you want this user to have:

   i. Enter `create`.

   ii. Enter `set scope=` followed by the scope name. Use tab-completion to see the list.

   iii. Enter `show` to see available filters, if any, and authorizations.

        iv. If a filter is available, set the filter value. Use tab-completion to see the list of possible filter values.

        v. Set to `true` all authorizations that you want this user to have.

        vi. Enter `commit`.

        Note that these authorizations can also be used to exclude authorizations that are granted to this user in a role. If you assign authorizations that have a more limited (more narrowly filtered) scope than the same authorizations that are granted in a role, then this user will only have the authorizations for the more limited scope.

    d. Enter `done`.

Alternative Method

To create a new user of the same type as an existing user and with the same roles and authorizations assigned, use the `clone` command. In `configuration users`, enter `clone` *`existing-user-name new-user-name`*. Set a `fullname` for the new user. If the type of the user that you are cloning is `local` or `data`, set a password.

Related Topics

- Understanding Users and Roles
- Managing User Properties
- User Authorizations
- Setting Appliance Preferences

Changing a User Password (BUI)

Use the following procedure to change a user's password. To change the password for any user other than yourself, you must have Super-User (root) privileges or a role with the user `changePassword` authorization.

1. From the **Configuration** menu, select **Users**.
2. Hover over the user in the **Users** list, and click the edit icon ✎ .
3. In the **Edit User** dialog box, type a new **Password**, and then type it again to confirm it.
4. Click **APPLY**.

Related Topics

- Editing Exceptions for a User (BUI)
- Editing Authorizations for a Role (BUI)

Changing a User Password (CLI)

Use the following procedure to change a user's password. To change the password for any user other than yourself, you must have Super-User (root) privileges or a role with the user `allow_changePassword` authorization.

1. Go to `configuration users`.
2. Enter `select` and the username of the user for which you want to change the password.
3. Set a new value for `initial_password`.

```
hostname:configuration users username> set initial_password=new password

                    initial_password = (set) (uncommitted)
```

4. Enter `commit`.

Related Topics

- Editing Exceptions for a User (CLI)
- Editing Authorizations for a Role (CLI)

Editing Exceptions for a User (BUI)

Use the following procedure to edit exceptions for a user.

1. From the **Configuration** menu, select **Users**.
2. Hover over the user in the **Users** list, and click the edit icon ✎ .
3. In the **Edit User** dialog box, click **Exceptions** to add or delete authorizations.
4. Optional: Add authorizations for this user.

   See "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

   Iterate the following steps until you have added all of the authorizations that you want this user to have:

   a. Select a **Scope**.

      Any filters that are available for this scope appear below the **Scope** selector.

   b. Specify filters for the scope as necessary.
   c. Click the check box for each authorization to add.
   d. Click **ADD** in the **Exceptions** section.

      The authorizations are listed at the bottom of the **Exceptions** section.

5. Optional: Delete authorizations for this user.

   In the list of authorizations at the bottom of the **Exceptions** section, hover over the authorization that you want to delete, and click the trash icon 🗑 .

6. Click **APPLY** at the top of the dialog box.

Related Topics

- Understanding Users and Roles
- User Authorizations
- Managing User Properties

Editing Exceptions for a User (CLI)

Use the following procedure to edit exceptions for a user.

1. Go to `configuration users`.
2. Enter `select` followed by the username.
3. Enter `exceptions` and then enter `show`.
4. Optional: Add authorizations for this user.

See "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

Iterate the following steps until you have added all of the authorizations that you want this user to have:

a. Enter `create`.

b. Enter `set scope=` followed by the scope name. Use tab-completion to see the list.

c. Enter `show` to see available filters, if any, and authorizations.

d. If a filter is available, set the filter value.

   Use tab-completion to see the list of possible filter values.

e. Set to `true` all authorizations that you want to add for this user.

f. Enter `commit`.

5. Optional: Delete authorizations for this user.

   For each authorization that you want to remove for this user, enter `destroy` and the name of the authorization.

```
hostname:configuration users username exceptions> destroy auth-001
This will destroy "auth-001". Are you sure? (Y/N) y
```

6. Enter `done` and then enter `done` again.

Related Topics

- Understanding Users and Roles
- User Authorizations
- Managing User Properties

Adding a Role (BUI)

A role is a collection of authorizations that can be assigned to a user. Use this procedure to define a new role. Also see the alternative method at the end of this task, which is ideal for cloning a local role as a directory role.

1. From the **Configuration** menu, select **Users**.

2. Click the add icon ⊕ next to **Roles**.

   See also "Alternative Method" following this procedure.

3. In the **Add Role** dialog box, set the role type, name of the role, and provide a description. Role types:

   - **Local** - Role applies to this appliance only.
   - **Directory** - Role applies to one of two directory group types and allows logging in as an administrator:
     - **LDAP** - Role applies to same-named, existing LDAP directory group. For **Name**, enter the exact same name for the LDAP directory group as configured on the LDAP server. Members of the same-named UNIX group are assigned this role and can log in as an administrator.
     - **Active Directory** - Role applies to same-named, existing Active Directory (AD) group. For **Name**, enter the exact same name in the format *name@domain* as configured for the AD group members on the AD server. Valid members of the same-named AD group are assigned this role and can log in as an administrator.

4. In the **Authorizations** section, add authorizations for this role.

   See "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

   Iterate the following steps until you have added all of the authorizations that you want this role to have:

   a. Select a **Scope**.

      Any filters that are available for this scope appear below the **Scope** selector.

   b. Specify filters for the scope as necessary.

   c. Click the check box for each authorization to add.

   d. Click **ADD** in the **Authorizations** section.

      The authorizations are listed at the bottom of the **Authorizations** section.

5. Click **ADD** at the top of the dialog box.

   The new role appears in the **Roles** list.

Alternative Method

To create a new role with the same authorizations as an existing role, hover over the entry for the existing role and click the clone icon ⊞ . Provide a role name, select **Local** or **Directory**, and click **ADD** at the top of the **Clone Role** dialog box. The new role type can be different from the existing type. For example, a local role can be cloned to a directory role.

Related Topics

- Understanding Users and Roles
- User Authorizations
- Managing User Properties

Adding a Role (CLI)

A role is a collection of authorizations that can be assigned to a user. Use this procedure to define a new role. Also see the alternative method at the end of this task, which is ideal for cloning a local role as a directory role.

1. Go to `configuration roles`.

2. Enter either `local` or `directory`, followed by the name of the role that you want to create. Role types:

   - `local` - Role applies to this appliance only.
   - `directory` - Role applies to one of two directory group types and allows logging in as an administrator:
     - **LDAP** - Role applies to same-named, existing LDAP directory group. For *name*, enter the exact same name for the LDAP directory group as configured on the LDAP server. Members of the same-named UNIX group are assigned this role and can log in as an administrator.
     - **Active Directory** - Role applies to same-named, existing Active Directory (AD) group. For *name*, enter the exact same name in the format *name@domain* as configured for the AD group members on the AD server. Valid members of the same-named AD group are assigned this role and can log in as an administrator.

3. Set the description of the role.

4. Enter `commit` to add the role.
5. Select the new role.
6. Enter `authorizations`.
7. Add authorizations for this role.

   See "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

   Iterate the following steps until you have added all of the authorizations that you want this role to have:

   a. Enter `create`.

   b. Enter `set scope=` followed by the scope name. Use tab-completion to see the list.

   c. Enter `show` to see available filters, if any, and authorizations.

   d. If a filter is available, set the filter value.

      Use tab-completion to see the list of possible filter values.

   e. Set to `true` all authorizations that you want to include in this role.

   f. Enter `commit`.
8. Enter `done` and then enter `done` again.

Alternative Method

To create a new role with the same authorizations as an existing role, use the `clone` command. In `configuration roles`, enter `clone` *`existing-role-name new-role-name new-role-type`* . For *new-role-type*, enter `local` or `directory`. The *new-role type* can be different from the existing type. For example, a local role can be cloned to a directory role. If no role type is specified, the new type is the same as the cloned type.

Related Topics

- Understanding Users and Roles
- User Authorizations
- Managing User Properties

Editing Authorizations for a Role (BUI)

A role is a collection of authorizations that can be assigned to a user. Use this procedure to add and delete authorizations for a role.

1. From the **Configuration** menu, select **Users**.
2. Hover over the role in the **Roles** list, and click the edit icon ✎ .
3. Optional: In the **Authorizations** section of the **Edit Role** dialog box, add authorizations for this role.

   See "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

   Iterate the following steps until you have added all of the authorizations that you want this role to have:

   a. Select a **Scope**.

      Any filters that are available for this scope appear below the **Scope** selector.

   b. Specify filters for the scope as necessary.

   c. Click the check box for each authorization to add.

   d. Click **ADD** in the **Authorizations** section.

      The authorizations are listed at the bottom of the **Authorizations** section.

4. Optional: Delete authorizations for this role.

   In the list of authorizations at the bottom of the **Authorizations** section, hover over the authorization that you want to delete, and click the trash icon 🗑 .

5. Click **APPLY** at the top of the dialog box.

Related Topics

- Understanding Users and Roles
- User Authorizations
- Managing User Properties

Editing Authorizations for a Role (CLI)

A role is a collection of authorizations that can be assigned to a user. Use this procedure to add and delete authorizations for a role.

1. Go to `configuration roles`.
2. Enter `select` followed by the role name.
3. Enter `authorizations`.
4. Optional: Add authorizations for this role.

   See "Scopes, Filters, and Authorizations Available for Users and Roles" in User Authorizations.

   Iterate the following steps until you have added all of the authorizations that you want this role to have:

   a. Enter `create`.

   b. Enter `set scope=` followed by the scope name. Use tab-completion to see the list.

   c. Enter `show` to see available filters, if any, and authorizations.

   d. If a filter is available, set the filter value.

      Use tab-completion to see the list of possible filter values.

   e. Set to `true` all authorizations that you want to include in this role.

   f. Enter `commit`.

5. Optional: Delete authorizations for this role.

   For each authorization that you want to remove for this role, enter `destroy` and the name of the authorization.

   ```
   hostname:configuration roles rolename authorizations> destroy auth-001
   This will destroy "auth-001". Are you sure? (Y/N) y
   ```

6. Enter `done` and then enter `done` again.

Related Topics

- Understanding Users and Roles

- User Authorizations
- Managing User Properties

Adding a User Who Can Only View the Dashboard (BUI)

Use the following procedure to add a user who can only view the dashboard.

1. From the **Configuration** menu, select **Users**.
2. Click the add icon ⊕ next to **Users**.
3. In the **Add User** dialog box, choose either **Directory** or **Local** for **Type** of user.
4. Enter values for **Username**, **Full Name**, and **Password**.

   The **Username** and **Password** properties are required. For help with setting the **Username** value, see Understanding Users and Roles.
5. Select the **Kiosk user** check box.
6. Ensure the **Kiosk screen** is set to status/dashboard.
7. Click **ADD**.

Related Topics

- Understanding Users and Roles
- Managing User Properties

Adding a User Who Can Only View the Dashboard (CLI)

Use the following procedure to add a user who can only view the dashboard.

1. Go to `configuration users`.
2. Enter either `directory` or `local` user type followed by a username.

   For descriptions of user types and for help with setting the username value, see Understanding Users and Roles.
3. Set a value for `initial_password`.
4. Enter `commit`.
5. Select the new user.
6. Set `kiosk_mode` to `true`.
7. Ensure that the value of `kiosk_screen` is `status/dashboard`.
8. Enter `commit`.

Related Topics

- Understanding Users and Roles
- Managing User Properties

Determining the Current Logged-in User

To view the identity of the current logged-in user, see the name at the top of the dashboard in the BUI, or use the `whoami` command in the CLI.

1. In the BUI, the full name of the user is shown to the left of the **Logout** button at the top of the dashboard.

   ```
   Full Name@hostname
   ```

   If a full name was not provided for the user, then no user identification is shown:

   ```
   hostname
   ```

   The full name might also appear in the tab of your browser.

2. In the CLI, the `whoami` command returns the login name of the user.

   ```
   hostname:> whoami
   loginname
   ```

Understanding Users and Roles

A user is one of the types shown in the following two tables. Only administrator types can be assigned authorizations or roles.

Table 2-31 Administrator User Types

BUI User Type	CLI User Type	Description
Local	`local`	- This appliance administrator is defined for this appliance only. - The username must be a new UNIX username. - A custom UID can be specified; otherwise, the system will assign the UID. - A password must be specified. - This user can be granted authorizations directly or by assigning custom roles. - Although local users are supported for data services, local groups are not supported.
Directory	`directory`	- This appliance administrator is managed by a directory service: NIS, LDAP, or Active Directory (AD). See NIS Configuration, LDAP Configuration, or Active Directory Configuration. - The user must be an existing UNIX NIS/LDAP user or an AD *name@domain* user. - User ID and Password are automatically assigned and cannot be set. If both NIS and LDAP are configured on the appliance and the services return different information for a particular user, the appliance uses the data provided by NIS. - When the appliance RADIUS service is enabled, *all* directory users log in using RADIUS. - This user can be granted authorizations directly or by assigning custom roles.
Auto	`auto`	This user type is automatically created when a user belonging to a directory role, but who was not explicitly added, logs in to the appliance for the first time. This then allows the user to set preferences, such as for the initial login screen and the session timeout duration. For more information about configuring user preferences, see Setting Preferences - BUI, CLI.

Table 2-32 Non-Administrator User Types

BUI User Type	CLI User Type	Description
Data-only	`data`	• A data-only user is defined locally for data (such as SMB, NFS, FTP) with no administrator access. • The username must be a new UNIX username. • A custom UID can be specified; otherwise, the system will assign the UID. • A password must be specified.
No-login	`nologin`	• A no-login user is not allowed to log in to the appliance. A username and UID are reserved for identity mapping purposes. • The username must be a new UNIX username. • A custom UID can be specified; otherwise, the system will assign the UID.

A role is a collection of authorizations that can be assigned to an administrator user type. Administrator users are assigned the "basic" role by default. The basic role enables the user to log in to the administrative interface and read most system configuration parameters. The basic role does not allow a user to make changes to the system. A user can be assigned additional roles and can be assigned additional authorizations directly. A role can be edited to add or delete authorizations.

Using roles is more secure than giving users the root password.

- Use roles to easily grant users only the set of authorizations that they require. For example, different roles could have authorizations to modify different services.
- Because users are operating under their own user names, you can more easily identify which real person performed a particular action.

A directory role specifically associates a role with an existing LDAP group or Active Directory (AD) group with the same name. As an example for LDAP, role "ZFS_Admins" is associated with LDAP group "ZFS_Admins". By creating the same LDAP directory role on multiple appliances, administrative privileges are granted to members of that LDAP group. Add or remove LDAP group members on the LDAP server configured for the appliances to centrally control who can log in to the appliance as an administrator. Also, on each appliance, you can assign different authorizations for the same directory role.

An automatic directory user is created when a user belonging to a directory role, but who was not explicitly added, logs in to the appliance for the first time. When automatic directory users are no longer authorized to be administrators, remove multiple users at once by using workflow "Destroy Unauthorized Directory Users" or remove them individually by manually removing them in the configuration-users area of the appliance software. For information about executing workflows, see Uploading and Executing Workflows Using the BUI and Executing Workflows using the CLI.

Related Topics

- Managing User Properties
- User Authorizations
- Adding an Administrator or User - BUI, CLI
- Changing a User Password - BUI, CLI
- Editing Exceptions for a User - BUI, CLI

- Adding a Role - BUI, CLI
- Editing Authorizations for a Role - BUI, CLI
- Adding a User Who can Only View the Dashboard - BUI, CLI

User Authorizations

Authorizations allow users to perform specific tasks, such as creating shares, rebooting the appliance, and updating the system software.

Authorizations are grouped into scopes. A particular scope might have a set of filters to narrow the scope of the authorization. For example, rather than an authorization to restart all services, a filter can be used so that this authorization can restart only the HTTP service.

The following table shows the available authorizations.

Table 2-33 Scopes, Filters, and Authorizations Available for Users and Roles

Scope BUI	Scope CLI	Filters	BUI/CLI Authorizations
Active Directory	`ad`	Domain or workgroup name	• **domain**/`allow_domain`: Join an Active Directory domain • **workgroup**/`allow_workgroup`: Join a workgroup
Alerts	`alert`	-	• **configure**/`allow_configure`: Create custom alert actions or threshold alerts • **post**/`allow_post`: Post custom alerts
Analytics	`stat`	List of drilldowns	• **configure**/`allow_configure`: Configure analytics hostname lookup policy • **create**/`allow_create`: Create a statistic with this drilldown present • **read**/`allow_read`: Read a statistic with this drilldown present

Table 2-33 (Cont.) Scopes, Filters, and Authorizations Available for Users and Roles

Scope BUI	Scope CLI	Filters	BUI/CLI Authorizations
Appliance	`appliance`	Appliance name	- **audit**/`allow_audit`: Emit an audit log entry - **configBackup**/`configBackup`: Create a configuration backup. Because data included in a configuration backup might be sensitive and because the `configBackup` authorization has the same full privileges as for the `root` user, be sure to read Security Considerations for Configuration Backups in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*. - **configExport**/`configExport`: Export a saved configuration. Because data included in a configuration backup might be sensitive, be sure to read Security Considerations for Configuration Backups in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*. - **configImport**/`configImport`: Import a saved configuration. Because data included in a configuration backup might be sensitive, be sure to read Security Considerations for Configuration Backups in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*. - **configRestore**/`configRestore`: Restore a saved configuration. Because data included in a configuration backup might be sensitive and because the `configRestore` authorization has the same full privileges as for the `root` user, be sure to read Security Considerations for Configuration Backups in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*. - **factoryReset**/`allow_factoryReset`: Restore the appliance to factory defaults - **notification-suspend**/`allow_notification-suspend`: Suspend all notifications - **peerSetup**/`allow_peerSetup`: Set up replication relations - **powerOff**/`allow_powerOff`: Power down the appliance

Table 2-33 (Cont.) Scopes, Filters, and Authorizations Available for Users and Roles

Scope BUI	Scope CLI	Filters	BUI/CLI Authorizations
			• **reboot**/`allow_reboot`: Reboot the appliance • **setName**/`allow_setName`: Modify the appliance name • **setTime**/`allow_setTime`: Set the appliance time • **shell**/`allow_shell`: Access the underlying Oracle Solaris shell • **systemCert**/`allow_systemCert`: Configure system certificates • **trustedCert**/`allow_trustedCert`: Configure trusted certificates
Cloud targets	`cloud`	Cloud target name	• **backup**/`allow_backup`: Backup snapshot data to cloud • **delete**/`allow_delete`: Delete cloud snapshots • **restore**/`allow_restore`: Restore cloud snapshots to local shares
Clustering	`cluster`	-	• **failback**/`allow_failback`: Failback resources to a cluster peer • **linkReset**/`allow_linkReset`: Reset a failed cluster I/O device • **takeover**/`allow_takeover`: Takeover resources from a cluster peer • **transfer**/`allow_transfer`: Transfer resources to a cluster peer
Datasets	`dataset`	-	• **configure**/`allow_configure`: Configure dataset retention policies and dataset auto suspend policies
Hardware	`hardware`	-	• **disk**/`allow_disk`: Online and offline disks • **disk-fault**/`allow_disk-fault`: Manually fault a disk • **led**/`allow_led`: Configure LEDs on disks, appliance, and external enclosures • **serviceProcessor**/`allow_serviceProcessor`: Configure network properties for the service processor • **storage-cancelSpare**/`allow_storage-cancelSpare`: Remove a drive as a hot spare • **storage-configure**/`allow_storage-configure`: Configure a storage pool • **storage-unconfigure**/`allow_storage-unconfigure`: Unconfigure a storage pool

Table 2-33 (Cont.) Scopes, Filters, and Authorizations Available for Users and Roles

Scope BUI	Scope CLI	Filters	BUI/CLI Authorizations
Keystores	`keystore`	Keystore name	• **listKeystore**/ `allow_listKeystore`: List keys present in a per-user keystore • **modifyKeystore**/ `allow_modifyKeystore`: Permit keystore modifications • **readKeystore**/ `allow_readKeystore`: Permit read access to sensitive values in a keystore
Networking	`net`	-	• **configure**/`allow_configure`: Configure networking devices, datalinks, and interfaces

Table 2-33 (Cont.) Scopes, Filters, and Authorizations Available for Users and Roles

Scope BUI	Scope CLI	Filters	BUI/CLI Authorizations
Projects and shares	`nas`	• Storage pool • Project • Share	• **backup**/`allow_backup`: Backup share data • **changeAccessProps**/`allow_changeAccessProps`: Configure who can access a share • **changeGeneralProps**/`allow_changeGeneralProps`: Change general properties on a share • **changeProtocolProps**/`allow_changeProtocolProps`: Configure protocol-specific properties • **changeSpaceProps**/`allow_changeSpaceProps`: Change quota and reservation on a share • **changeUserQuota**/`allow_changeUserQuota`: Change user and group quotas on a share • **clearLocks**/`allow_clearLocks`: Clear locks held on behalf of an NFS client • **clone**/`allow_clone`: Clone a snapshot to a normal filesystem • **createProject**/`allow_createProject`: Create a project • **createShare**/`allow_createShare`: Create a filesystem or LUN • **destroy**/`allow_destroy`: Remove a project or share • **destroySnap**/`allow_destroySnap`: Remove a snapshot • **disableLockedSnap**/`allow_disableLockedSnap`: Disable retention on schedule of snapshots • **enableLockedSnap**/`allow_enableLockedSnap`: Enable and configure retention on schedule of snapshots • **encryption**/`allow_encryption`: Manage encryption keys for a pool, project, or share • **promote**/`allow_promote`: Promote a clone • **releaseSnapRetention**/`allow_releaseSnapRetention`: Release a snapshot retention policy

Table 2-33 (Cont.) Scopes, Filters, and Authorizations Available for Users and Roles

Scope BUI	Scope CLI	Filters	BUI/CLI Authorizations
			• **rename**/`allow_rename`: Rename a project or share • **renameSnap**/`allow_renameSnap`: Rename a snapshot • **restore**/`allow_restore`: Restore data to share • **retainSnap**/`allow_retainSnap`: Retain a snapshot • **retentionAuto**/ `allow_retentionAuto`: Enable automatic retention • **retentionMandatory**/ `allow_retentionMandatory`: Enable mandatory retention • **retentionPeriods**/ `allow_retentionPeriods`: Alter retention periods • **rollback**/`allow_rollback`: Rollback a filesystem to a previous snapshot • **rrsource**/`allow_rrsource`: Configure data replication to other appliances • **rrtarget**/`allow_rrtarget`: Manage data replicated from other appliances • **scheduleLockedSnap**/ `allow_scheduleLockedSnap`: Configure retention on schedule of snapshots • **scheduleSnap**/ `allow_scheduleSnap`: Configure a recurring schedule of snapshots • **scrub**/`allow_scrub`: Check a storage pool for errors • **shadowMigration**/ `allow_shadowMigration`: Manage shadow migration on a share • **takeSnap**/`allow_takeSnap`: Take a manual snapshot
Roles	`role`	Role name	• **changeAuths**/ `allow_changeAuths`: Configure authorizations for a role • **changeDescription**/ `allow_changeDescription`: Change a description of a role • **create**/`allow_create`: Create a role • **destroy**/`allow_destroy`: Destroy a role

Table 2-33 (Cont.) Scopes, Filters, and Authorizations Available for Users and Roles

Scope BUI	Scope CLI	Filters	BUI/CLI Authorizations
SAN	`stmf`	-	• **configure**/`allow_configure`: Configure SAN hosts and targets
Services	`svc`	Service name	• **administer**/`allow_administer`: Enable or disable service • **configure**/`allow_configure`: Configure service properties and settings • **restart**/`allow_restart`: Restart service
Shares property schema	`schema`	-	• **modify**/`allow_modify`: Modify property schema
Update	`update`	-	• **delete**/`allow_delete`: Delete system updates • **update**/`allow_update`: Update system software • **upload**/`allow_upload`: Upload system updates
Users	`user`	Username	• **changeAuths**/ `allow_changeAuths`: Configure authorizations for a user • **changePassword**/ `allow_changePassword`: Change a password • **changePreferences**/ `allow_changePreferences`: Configure preferences for a user • **changeProperties**/ `allow_changeProperties`: Configure properties for a user • **changeRESTTokens**/ `allow_changeRESTTokens`: Configure REST tokens for a user • **changeRoles**/ `allow_changeRoles`: Configure roles for a user • **create**/`allow_create`:Create a user • **destroy**/`allow_destroy`: Destroy a user
Workflow	`workflow`	• Owner • Name	• **modify**/`allow_modify`: Delete workflow • **read**/`allow_read`: Execute workflow
Worksheet	`worksheet`	• Owner • Name	• **modify**/`allow_modify`: Modify worksheet • **read**/`allow_read`: Read worksheet

Related Topics

- Adding an Administrator or User - BUI, CLI

- Editing Exceptions for a User - BUI, CLI
- Adding a Role - BUI, CLI
- Editing Authorizations for a Role - BUI, CLI

Managing User Properties

The BUI **Configuration: Users** page lists both users and roles. Hover over an entry to see the clone, edit, and destroy buttons for the user or role. Double-click a user or role, or click its edit icon ✎ to display its **Edit User** or **Edit Role** dialog box. The following table describes the buttons on this page.

Table 2-34 BUI Users Page Buttons

Button	Description
⊕	Add a new user or role. A new **Add User** or **Add Role** dialog box is displayed where you enter values for properties.
🔍	Open a search box. Enter a search string to display only user or role entries in which that search text appears in the listing. Click the search button again, or click **Show All**, at the top of the list, to display the full list.
⊞	Clone a user or role. Add a new user or role with the same authorizations as the cloned user or role. Specify a role type for the clone, as defined in Role Properties. If no type is specified, the clone will have the same type as the original role.
✎	Edit a user or role.
🗑	Remove a user, role, or authorization.

Depending on the type of user, all of the following properties can be set when adding a user. A subset of these properties can be set when editing a user.

Table 2-35 User Properties

BUI Property	CLI Property	Description
Type	`type`	For a description of user types, see Understanding Users and Roles.
Username	`logname`	Unique login name for the user.
User ID	`uid`	Enabled only for **Local**, **Data**, and **No-login** users. You can specify the user ID or allow the system to assign the user ID. If you specify the user ID, the user ID cannot be less than 100, cannot be greater than 2147483646, and cannot be equal to 60001, 60002, or 65534. Those UIDs are reserved by the operating system vendor for use in future applications. Their use by end-system users or vendors of layered products is not supported and can cause security issues with other applications.
Full Name	`fullname`	Full name or real name for the user. In the BUI, the full name is shown to the left of the **Logout** button at the top of the dashboard, and it might also be shown on the browser tab.

Table 2-35 (Cont.) User Properties

BUI Property	CLI Property	Description
Password/Confirm	`initial_password`	For **Local** and **Data** users, type the initial password in these fields.
Require session annotation	`require_annotation`	When enabled: - **BUI** – Require the user to enter a comment prior to displaying the initial BUI page. - **CLI** – Require the user to enter a comment prior to displaying the CLI prompt. - **REST** – Requests fail as unauthorized. The comment appears in the audit log. This annotation can be used to describe the purpose of the login. A ticket number could be used to track particular project work.
Kiosk user	`kiosk_mode`	When enabled, this user is a kiosk user: - **BUI** – The user is restricted to viewing only the screen that is the value of the `Kiosk screen` property. - **CLI** – Login fails. - **REST** – Requests fail as unauthorized.
Kiosk screen	`kiosk_screen`	The screen that this user is restricted to if **Kiosk** user is enabled or `kiosk_mode` is true. Default: `status/dashboard`
Roles	`roles`	The roles assigned to a directory or local user.
Group Assigned	`group_roles`	Implicit role assignment based on user's group membership, for both directory and automatic directory users. This role cannot be modified.
Exceptions	`exceptions`	Additional authorizations assigned to a directory or local user, or limitations on authorizations that are assigned in a role.
-	`preferences`	User environment preferences such as locale, BUI start page, timeouts, SSH public keys, and REST login tokens. See Setting Appliance Preferences.

The following properties can be set when creating or editing a role.

Table 2-36 Role Properties

BUI Property	CLI Property	Description
Name	`name`	Name of the role as it will be shown in lists.
Description	`description`	Verbose description of the role.
Authorizations	`authorizations`	Authorizations for this role.

Table 2-36 (Cont.) Role Properties

BUI Property	CLI Property	Description
Type	`type`	Type of role: • **Local** - Applies to this appliance only. • **Directory** - Applies to the appliances controlled by a directory service: NIS, LDAP, or Active Directory (AD).

Related Topics

- Adding an Administrator or User - BUI, CLI
- Changing a User Password - BUI, CLI
- Editing Exceptions for a User - BUI, CLI
- Adding a Role - BUI, CLI
- Editing Authorizations for a Role - BUI, CLI
- Adding a User Who can Only View the Dashboard - BUI, CLI

Setting Appliance Preferences

This section describes how to configure user preferences such as locality, CLI session properties, advanced analytics, and SSH keys. The CLI allows you to set preferences for any user for which you have the required authorization. The HTTPS service controls the session timeout for the BUI, and has a default value of 15 minutes. For more information, see HTTPS Properties and Logs.

To configure user preferences, use the following sections:

- Setting Preferences - BUI, CLI
- Setting SSH Public Keys - BUI, CLI
- Preference Properties

Setting Preferences (BUI)

Use this procedure to set preferences for the user that you are currently logged in as.

To change preferences for a user other than the user that you are currently logged in as, see Setting Preferences (CLI).

1. From the **Configuration** menu, select **Preferences**.
2. Set preferences values.

   See Preference Properties for descriptions.

   By default, there is no limit on the length of time that the CLI can be idle (`infinite`). To set a limit on the length of time that the CLI can be idle, enter a positive integer for CLI idle timeout, and select the time units from the menu. If the timeout limit is reached, the CLI is closed.

   To set the BUI session timeout, see HTTPS Properties and Logs.

   For more information about SSH public keys, see Setting SSH Public Keys (BUI).

3. Click **APPLY**.

Related Topics

- Adding an Administrator or User (BUI)
- Setting Preferences (CLI)
- Preference Properties
- HTTPS Properties and Logs
- Setting SSH Public Keys (BUI)

Setting Preferences (CLI)

Use this procedure to set preferences for the user that you are currently logged in as or for any user for which you have the changePreferences authorization. See User Authorizations and "Editing Exceptions for a User" - BUI, CLI for information about gaining the changePreferences authorization.

1. Go to configuration users.

   ```
   hostname:> configuration users
   ```

2. Select the user for whom you want to edit preferences.

   a. Enter show to list users.

      To edit preferences for a user other than the user that you are currently logged in as, you must have the changePreferences authorization for that user.

   b. Use the name in the USERNAME column:

      ```
      hostname:configuration users> select username
      ```

3. Enter preferences, and then enter show to list the preferences.

   Note that the session_annotation property only shows for the currently logged in user.

   ```
   hostname:configuration users username> preferences
   hostname:configuration users username preferences> show
   Properties:
                           locale = C
                     login_screen = status/dashboard
                 cli_idle_timeout = infinite (default)
               session_annotation =
               advanced_analytics = false

   Children:
                                keys => Manage SSH public keys
                              tokens => Manage REST tokens
   ```

4. Set preferences values.

   See Preference Properties for descriptions.

   By default, there is no limit on the length of time that the CLI can be idle (infinite). To set a limit on the length of time that the CLI can be idle, set a value such as 15minutes or 1hour for cli_idle_timeout. The default time units are seconds, so if you enter a value of 1800 (with no units specified), the timeout value is 1800 seconds or 30 minutes. If the timeout limit is reached, the CLI is closed. To disable the timeout, set the value of cli_idle_timeout to infinite or specify the following command:

   ```
   hostname:configuration preferences> unset cli_idle_timeout
                 cli_idle_timeout = infinite (default, uncommitted)
   ```

For more information about keys, see Setting SSH Public Keys (CLI).

5. Enter `commit`.
6. Enter `done`.
7. Optional: Alternatively, use `configuration preferences`.

   This method only allows you to set preferences for the user that you are currently logged in as.

   a. Go to `configuration preferences`.
   b. Enter `show`.
   c. Set preferences values, enter `commit`, enter `done`.

Related Topics

- Adding an Administrator or User (CLI)
- Preference Properties
- Setting SSH Public Keys (CLI)

Setting SSH Public Keys (BUI)

SSH public keys can be used to allow SSH connections without the use of passwords. This feature is useful for administrator convenience and for automated execution of scripts.

Use this procedure to set SSH public keys for the user that you are currently logged in as. To set keys for a user other than the user that you are currently logged in as, see Setting SSH Public Keys (CLI).

1. From the **Configuration** menu, select **Preferences**.
2. Click the add icon ⊕ next to **SSH Public Keys**.
3. Select a **Type**, and then type the SSH public key, and an optional key comment.
4. Click **ADD**.

Related Topics

- Setting Preferences (BUI)
- Setting SSH Public Keys (CLI)
- Preference Properties
- SSH Configuration

Setting SSH Public Keys (CLI)

SSH public keys can be used to allow SSH connections without the use of passwords. This feature is useful for administrator convenience and for automated execution of scripts.

Use this procedure to set SSH public keys for the user that you are currently logged in as or for any user for which you have the `changePreferences` authorization. See User Authorizations and "Editing Exceptions for a User" - BUI, CLI for information about gaining the `changePreferences` authorization.

1. Go to `configuration users`.
2. Select the user for whom you want to edit preferences.

```
hostname:configuration users> select username
```

3. Enter `preferences keys`, and then enter `show` to list any existing SSH keys.

   ```
   hostname:configuration users username> preferences keys
   hostname:configuration users username preferences keys> show
   ```

4. Configure an SSH key.
   1. Enter `create` to create a new SSH key.
   2. Set the `type` (`DSA` or `RSA`), the `key`, and optionally a `comment` that gives the purpose of this key.
5. Enter `commit`.

   Enter `show` to view the new key.
6. Enter `done`.
7. Optional: Alternatively, use `configuration preferences`.

   This method only allows you to create SSH keys for the user that you are currently logged in as.
   1. Go to `configuration preferences keys`.
   2. Enter `create`.
   3. Set the `type`, `key`, and optionally `comment`.
   4. Enter `commit`, enter `show`, and enter `done`.

Related Topics

- Setting Preferences (CLI)
- Preference Properties
- SSH Configuration

Preference Properties

The following table describes the properties for setting user preferences.

Table 2-37 Preference Properties

BUI Property	CLI Property	Description
Initial login screen	`login_screen`	The BUI page that is presented upon successful login if a page is not specified in the URL. By default, this page is the **Status Dashboard** (`status/dashboard`). The screen that is presented prior to login includes the system login prompt and an optional system message. See “System Login Message” in System Identity Properties and Logs.

Table 2-37 (Cont.) Preference Properties

BUI Property	CLI Property	Description
Locality	`locale`	C by default. C and POSIX localities support only ASCII characters or plain text. ISO 8859-1 supports the following languages: Afrikaans, Basque, Catalan, Danish, Dutch, English, Faeroese, Finnish, French, Galician, German, Icelandic, Irish, Italian, Norwegian, Portuguese, Spanish, and Swedish.
CLI idle timeout	`cli_idle_timeout`	The length of time that the CLI can be idle before the session is closed. The default value, `infinite`, means the CLI will not automatically close when idle.
Current session annotation	`session_annotation`	Annotation text added to audit logs.
Advanced analytics statistics	`advanced_analytics`	Make available additional statistics in Analytics.
SSH Public Keys	`keys`	RSA/DSA public keys. Text comments can be associated with the keys to help administrators track why they were added.
REST Tokens	`tokens`	Persistent and non-persistent REST login tokens. Set the token `name`, set `preserve` to `true` or `false`, and set the token `expiration` in seconds.

Configuring Alerts

Important Oracle ZFS Storage Appliance events, such as hardware and software faults, trigger alerts. Alerts appear in the alert logs.

This section describes how to do the following tasks:

- Adding Alert Actions
  - Configure additional alert actions (responses to event alerts) for specified system events.
  - Configure alert actions for events that are defined in a workflow. In addition to the same alert actions that you can specify for system events, these alert actions allow you to provide information that is provided by the system for system events, such as what effect this event has on the system and the recommended action for the administrator to take.
  - Add an alert action for threshold alerts for specific statistic events.
- Adding Threshold Alerts

  Create threshold alerts for specific Analytics statistics.

Adding Alert Actions

An alert action is a response to an event alert. The procedures in this section describe how to select one or more events and specify responses to alerts for those events. Examples of alert responses include sending an email, resuming a dataset, or executing a workflow. More than one alert action can be specified for any particular event alert.

To create alert actions, use the following topics:

- Adding an Alert Action - BUI, CLI
- Adding an Alert Action for an Event Defined in a Workflow - BUI, CLI
- Adding an Action for a Specific Threshold Alert - BUI, CLI
- Alert Event Categories
- Alert Action Types

Adding an Alert Action (BUI)

Select one or more events and select responses to alerts for those events. More than one alert action can be specified for a particular event alert.

1. From the **Configuration** menu, select **Alerts**.
2. Click the add icon ⊕ in the **Alert actions** tab title.
3. In the **Events** section of the dialog box, select the event **Category**.

   Except for **Custom**, events within the category are displayed when you select the category. Select one of the following:

   - **All events** - Select either All (All events) or Subset. If you select Subset, toggle the checkboxes for the events so that only the types of events (such as service alerts or hardware faults) for which you want to perform this alert action are checked.
   - **A specific category, such as Hardware events or Services** - Select either All (All events in this category) or Subset. If you select Subset, toggle the checkboxes for the events so that only the events for which you want to perform this alert action are checked.
     The Analytics category enables you to select high-level events such as datasets auto-suspend warning, memory total exceeded, and usage exceeded. To define an alert for a threshold condition for a specific statistic, see Adding a Threshold Alert (BUI).

     The Thresholds category does not enable you to specify threshold events for a specific statistic. The Thresholds category enables you to select high-level events such as statistic threshold violated. Configuring alert actions for these high-level threshold events might be adequate so that you do not need to configure a separate alert action for each specific statistic threshold event alert that you define. To define an alert action that is specific to a particular statistic threshold event, add a threshold alert action as described in Adding an Action for a Specific Threshold Alert (BUI).
   - **Custom** - This selection enables you to specify an alert action for an event that you define in a workflow. See Adding an Alert Action for an Event Defined in a Workflow (BUI).
4. In the **Alert actions** section of the dialog box, select the action to take when this event alert is sent.

   Most of the actions have arguments. For example, you might have to specify an email recipient or select a dataset or workflow.

5. Optional: Select the **TEST** button to create a test alert, and execute this alert action.

   A test can be useful, for example, for checking whether email or SNMP is configured correctly.

6. Optional: To specify additional actions for this event alert, click the add icon ⊕ in the **Alert actions** section title. Select the action and specify arguments.
7. Click **ADD** at the top right of the dialog box.

Next Steps

- To modify an alert action, double-click the alert action or hover over the alert action and click the edit icon ✎ and make the changes. Then click **APPLY**.
- To delete an alert action, hover over the alert action, click the delete icon 🗑 and confirm that you want to delete this alert action.

Related Topics

- Alert Event Categories
- Alert Action Types
- Adding a Threshold Alert (BUI)
- Adding an Action for a Specific Threshold Alert (BUI)
- Creating and Posting Custom Alerts from Within a Workflow

Adding an Alert Action (CLI)

Specify one or more events and specify responses to alerts for those events. More than one alert action can be specified for a particular event alert.

1. Go to `configuration alerts actions` and enter the `create` command.

   ```
   hostname:configuration alerts actions> create
   ```

2. Set the `category` property.

   See Alert Event Categories or enter `set category=` followed by a tab character to see the list of available event categories.

   Specify one of the following:

   - **The `all` category** - Enter `show` to see the list of events in this category. By default, all events are set to `true`, so the alert action will be performed for all events in the category. If this alert action should be performed for only a subset of events, change to `false` the events that should not cause this alert action to be performed.
   - **A specific category, such as `hardware_faults` or `smf`** - Enter `show` to see the list of events in this category. By default, all events are set to `true`, so the alert action will be performed for all events in the category. If the alert action should be performed for only a subset of events, change to `false` the events that should not cause this alert action to be performed.
     The `analytics` category enables you to select high-level events such as datasets auto-suspend warning, memory total exceeded, and usage exceeded. To define an alert for a threshold condition for a specific statistic, see Adding a Threshold Alert (CLI).

     The `thresholds` category does *not* enable you to specify threshold events. The `thresholds` category requires that you specify an existing threshold event alert. To define an alert for a threshold condition for a specific statistic, see Adding a Threshold

Alert (CLI). To define an alert action that is specific to a particular statistic threshold event, add a threshold alert action as described in Adding an Action for a Specific Threshold Alert (CLI). To define an alert action for high-level threshold events such as statistic threshold violated, use the Thresholds category in the BUI procedure Adding an Alert Action (BUI).

- **The `custom` category -** The `custom` category enables you to define an alert action for an event that you define in a workflow. See Adding an Alert Action for an Event Defined in a Workflow (CLI).

3. Enter `commit`.
4. Enter `list` to see the list of all configured alert actions.

   The new alert action should be at the bottom of the list and should have a name (`actions-### `) and a category, but no action or handler.
5. Select the new alert action.

   ```
   hostname:configuration alerts actions> select actions-001
   hostname:configuration alerts actions-001>
   ```
6. Set the alert handler. Enter `action` and then enter `get`.

   ```
   hostname:configuration alerts actions-001> action
   hostname:configuration alerts actions-001 action (uncommitted)> get
                          handler = email
                          address = (unset)
                          subject = (unset)
   ```

   The default handler is `email`. If you want a different handler, enter `set handler=` followed by a tab character to see the list of possible handlers.
7. Set values for any handler arguments.
   1. Enter `get` again to see the list of arguments, if any, for the specified handler.
   2. Most of the actions have arguments. For example, you might have to specify an email recipient, dataset, or workflow.
8. Enter `commit`.

   Enter `list` to confirm that the action is correct.
9. Enter `done`.

   Enter `list` to view the list of actions.

Next Steps

- To change which events are included in this alert action, enter `select actions-###`, set the events to `true` or `false`, and enter `commit`.
- To specify additional actions for this alert action, enter `select actions-###`, and repeat step 6 through step 9. Additional actions are shown on separate lines for the specified `actions-###`.
- To delete an alert action, enter `destroy actions-###` and enter `y` to confirm.

Related Topics

- Alert Event Categories
- Alert Action Types
- Adding a Threshold Alert (CLI)

- Adding an Action for a Specific Threshold Alert (CLI)
- Creating and Posting Custom Alerts from Within a Workflow

Adding an Alert Action for an Event Defined in a Workflow (BUI)

This alert action is performed when the `postalert` function is called in a workflow, and the UUID of this alert action is specified. For more information, see Creating and Posting Custom Alerts from Within a Workflow.

Before You Begin

To perform this procedure, you must have the alerts `configure` authorization.

1. From the **Configuration** menu, select **Alerts**.
2. Click the add icon ⊕ in the **Alert actions** tab title.
3. In the **Events** section of the dialog box, select the **Custom** event category.
4. Select the severity of the event that precipitated this alert action: **Minor**, **Major**, or **Critical**.
5. Enter a brief **Description** of the event that precipitated this alert action.
6. Optional: Enter a **Response**, **Impact**, and **Recommended action**.
   - **Response** - String that describes actions that will be performed by the system to mitigate the effects of this event.
   - **Impact** - String that describes the effect that this event has on the appliance.
   - **Recommended action** - String that describes actions that the administrator should take to mitigate the effects of this event.
7. In the **Alert actions** section of the dialog box, select the action to take when this event alert is sent.

   Most of the actions have arguments. For example, you might have to specify an email recipient or select a dataset or workflow.
8. Optional: To specify additional actions for this event alert, click the add icon ⊕ in the **Alert actions** section title. Select the action and specify arguments.
9. Click **ADD** at the top right of the dialog box.

Related Topics

- Adding an Alert Action (BUI)
- Alert Event Categories
- Alert Action Types

Adding an Alert Action for an Event Defined in a Workflow (CLI)

This alert action is performed when the `postalert` function is called in a workflow, and the UUID of this alert action is specified. For more information, see Creating and Posting Custom Alerts from Within a Workflow.

Before You Begin

To perform this procedure, you must have the alerts `allow_configure` authorization.

1. Go to `configuration alerts actions` and enter the `create` command.
2. Set the `category` property to `custom`.

3. Set the `severity` of the event that precipitated this alert action: `Critical`, `Major`, or `Minor`.
4. Set a brief `description` of the event that precipitated this alert action.
5. Optional: Set a `response`, `impact`, and `recommended_action`.
   - **`response`** - String that describes actions that will be performed by the system to mitigate the effects of this event.
   - **`impact`** - String that describes the effect that this event has on the appliance.
   - **`recommended_action`** - String that describes actions that the administrator should take to mitigate the effects of this event.
6. Enter `commit`.
7. Enter `list` to see the list of all configured alert actions.

   The new alert action should be at the bottom of the list, and should have a name (`actions-###`) and a category, but no action or handler.
8. Select the new alert action.
9. Set the alert handler. Enter `action`, and then enter `get`.

   The default handler is `email`. If you want a different handler, enter `set handler=` followed by a tab character to see the list of possible handlers.
10. Set values for any handler arguments.
    a. Enter `get` again to see the list of arguments, if any, for the specified handler.
    b. Most of the actions have arguments. For example, you might have to specify an email recipient, dataset, or workflow.
11. Enter `commit`.

    Enter `list` to confirm that the action is correct.
12. Enter `done`.

    Enter `list` to view the list of actions.

Related Topics

- Adding an Alert Action (CLI)
- Alert Event Categories
- Alert Action Types

Adding an Action for a Specific Threshold Alert (BUI)

This procedure describes how to add an alert action for threshold alerts for specific statistic events.

If the specific threshold alert does not require a unique response, you can use the **Thresholds** category in the procedure Adding an Alert Action (BUI) to specify alert actions that execute for high-level threshold events, such as a violated statistic threshold.

1. From the **Configuration** menu, select **Alerts**.
2. Click the **Threshold alerts** tab.

   If the threshold event alert for which you want to add an alert action does not already exist, use the procedure Adding a Threshold Alert (BUI) to add the threshold alert. You can add alert actions when you create the threshold alert, and then skip the remaining steps in this procedure.

3. Select the threshold alert for which you want to add an alert action.

   Double-click the alert or hover over the alert and click the edit icon ✎ .

4. In the **Alert actions** section of the dialog box, select the action to take when this threshold alert is sent.

   Most of the actions have arguments. For example, you might have to specify an email recipient, or select a dataset or workflow.

5. Optional: Click the **TEST** button to create a test alert, and to execute this alert action.

   For example, a test can be useful for checking whether email or SNMP is configured correctly.

6. Optional: To specify additional actions for this threshold event alert, click the add icon ⊕ in the **Alert actions** section title.

   Select the action and specify arguments.

7. Click **APPLY** at the top right of the dialog box.

Related Topics

Adding a Threshold Alert (BUI)

Adding an Action for a Specific Threshold Alert (CLI)

This procedure describes how to add an alert action for threshold alerts for specific statistic events. The `thresholds` category requires that you specify an existing threshold event alert.

If the specific threshold alert does not require a unique response, you can use the `thresholds` category in the procedure Adding an Alert Action (CLI) to specify alert actions that execute for high-level threshold events, such as a violated statistic threshold.

1. Get the threshold alert UUID.

   a. Go to `configuration alerts thresholds` and enter the `list` command.

      If the threshold event alert for which you want to add an alert action does not already exist, use the procedure Adding a Threshold Alert (CLI) to add the threshold alert.

   b. Enter s`elect threshold-###` for the threshold alert for which you want to add an alert action.

   c. Enter `get uuid`, copy that threshold alert UUID, and enter `done`.

2. Go to `configuration alerts actions` and enter the `list` command.

   Do the following if an alert action that has a UUID that matches the UUID from step 1 does *not* already exist.

   a. Enter `create`.

   b. Enter `set category=thresholds`.

   c. Enter `set thresholdid=`*`uuid`*, where *uuid* is the UUID that you copied in step 1.

   d. Enter `commit`.

   If an alert action that has a UUID that matches the UUID from step 1 already exists, you can specify an additional action for this alert.

3. Select the alert action that has a UUID that matches the UUID from step 1.

4. Set the alert handler. Enter `action`, and then enter `get`.

The default handler is `email`. If you want a different handler, enter `set handler=` followed by a tab character to see the list of possible handlers.

5. Set values for any handler arguments.

   Enter `get` again to see the list of arguments, if any, for the specified handler.

   Most of the actions have arguments. For example, you might have to specify an email recipient, dataset, or workflow.

6. Enter `commit`.

   Enter `list` to confirm that the action is correct.

7. Enter `done`.

   Enter `list` to view the list of actions.

Related Topics

Adding a Threshold Alert (CLI)

Alert Event Categories

The following table describes the alert event categories that are available. BUI event categories are listed in the **Category** menu in the **Event** section of the **Add action** dialog box. Events within each category are listed when you select the category.

CLI event categories are values of the `category` property and can be listed by entering `set category=` followed by a tab character. Events within each category are listed when you specify the category and then issue the `show` command.

Table 2-38   Alert Action Event Categories

BUI Category	CLI Category	Description
Active Directory	`ad`	Active Directory or SMB Kerberos client authentication degraded.
All events	`all`	High-level events such as all alerts or defects, service alerts, and hardware faults.
Analytics	`analytics`	High-level events such as datasets auto-suspend warning, memory total exceeded, and usage exceeded.
Appliance Software	`appliance_software`	Events that prevent software update or that result in kernel panic.
Cloud Snapshot	`cloud`	Cloud operations, including backup and restore of share snapshots to the cloud.
Cluster	`cluster`	Cluster events, including link failures and peer errors.
Custom	`custom`	An alert action for a user-defined event, which is specified in a workflow. See Creating and Posting Custom Alerts from Within a Workflow.
Hardware events	`hardware`	Appliance boot and hardware configuration changes.
Hardware faults	`hardware_faults`	Any hardware fault.

Table 2-38 (Cont.) Alert Action Event Categories

BUI Category	CLI Category	Description

NDMP operations
NDMP: backup only
NDMP: restore only
	`ndmp`
	`backup`
	`restore`
		NDMP TAR/DUMP backup and restore start and finish events.

Network	`network`	Network port, datalink, and IP interface events and failures.
Phone home	`scrk`	Support bundle upload events.

Remote replication
Remote replication: source only
Remote replication: target only
	`replication`
	`replication_source`
	`replication_target`
		Send and receive events and failures.

Services	`smf`	Software services failure events.
Shadow migration events	`shadow`	Migration errors or migration complete.
Thresholds	`thresholds`	In the BUI, **Thresholds** events are `Stat threshold error`, `Stat threshold normal`, and `Stat threshold violated`. In the CLI, `thresholds` enables you to add an action to an existing threshold event alert as described in Adding an Action for a Specific Threshold Alert (CLI).
ZFS pool	`zfs_pool`	Storage pool events, including scrub and hot space activation.

Alert Action Types

The following table describes the alert action types that are available. BUI action types are listed in a menu in the **Alert actions** section of the **Add action** dialog box.

CLI action types are values of the `handler` property, and can be listed by entering `set handler=` followed by a tab character.

Table 2-39 Alert Action Types

BUI Action Type	CLI Action Type	Action Type Description
Send email Enter values for Send to and Subject	`email` Set `address` and `subject`.	Sends an email with the specified subject to the specified recipients. To send to multiple individual recipients, separate email addresses with a comma and a space on one line. Use the SMTP service to configure how email is sent. • **BUI** - **Configuration: Services: SMTP** • **CLI** - `configuration services smtp`

Table 2-39 (Cont.) Alert Action Types

BUI Action Type	CLI Action Type	Action Type Description
Send SNMP trap	`snmp_trap`	Sends an SNMP trap that contains alert details. Use the SNMP service to configure an SNMP trap destination. - **BUI** - **Configuration: Services: SNMP** - **CLI** - `configuration services snmp`
Send Syslog Message	`syslog`	Sends a system message that contains alert details to one or more remote systems. Use the **Syslog** service to configure syslog destinations. - **BUI** - **Configuration: Services: Syslog** - **CLI** - `configuration services syslog` For more information about sending syslog messages, see Syslog Configuration.
Resume dataset Select the dataset from the menu.	`resume_dataset` Set `dataset`.	Resumes an Analytics dataset. Resuming and suspending datasets can be useful for diagnosing intermittent performance issues, and for other cases when keeping a dataset continuously enabled is not desirable. For more information, see About Analytics Datasets in *Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x*.
Suspend dataset Select the dataset from the menu.	`suspend_dataset` Set `dataset`.	Suspends an Analytics dataset.
Resume worksheet Select the worksheet from the menu.	`resume_worksheet` Set `worksheet`.	Resumes an Analytics worksheet. Resuming and suspending worksheets can be useful for the same reasons as resuming and suspending datasets. A worksheet might contain numerous datasets. For more information, see Worksheet Graphs and Plots in *Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x*.
Suspend worksheet Select the worksheet from the menu.	`suspend_worksheet` Set `worksheet`.	Suspends an Analytics worksheet.
Execute workflow Select the workflow from the menu.	`execute_workflow` Set `workflow`.	Executes the specified workflow. To enable a workflow be eligible as an alert action, the `alert` action of the workflow must be set to `true` as described in Using Workflows for Alert Actions. **Note:** The executed workflow cannot post an alert. See Creating and Posting Custom Alerts from Within a Workflow.

Adding Threshold Alerts

A threshold alert is a custom alert in which a threshold is defined for a particular Analytics statistic, and the alert action is executed when the statistic value is outside that threshold. See also Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x.

To create custom Analytics statistics threshold alerts, see the following topics:

- Adding a Threshold Alert - BUI, CLI
- Threshold Alert Properties

Adding a Threshold Alert (BUI)

Use this procedure to specify an alert for a specific Analytics statistic threshold event, and to specify an action (response) for the alert.

Before You Begin

To perform this procedure, you must have the alerts `configure` authorization.

1. From the **Configuration** menu, select **Alerts**.
2. Click the **Threshold alerts** tab.
3. Click the add icon ⊕ in the **Threshold alerts** tab title.
4. In the **Threshold** section of the dialog box, specify the threshold event.
   a. Select a **statistic** from the drop-down menu.
   b. Select the **event type**: `exceeds` or `falls below`.
   c. Specify the percent or the number of bytes, operations, accesses, or requests.
5. In the **Timing** section of the dialog box, specify when to send this threshold event alert.

   You can specify to send the alert only when the event occurs on particular days or within particular time ranges or after the event has occurred for a particular length of time. You can specify a multiple of these timing parameters.
6. In the **Alert actions** section of the dialog box, select the action to take when this threshold event alert is sent.

   Most of the actions have arguments. For example, you might have to specify an email recipient, or select a dataset or workflow.

   If this threshold alert does not require a unique response, you can leave this **Alert actions** section blank, and instead use the **Thresholds** category in the procedure Adding an Alert Action (BUI) to specify alert actions that execute for high-level threshold events, such as a violated statistic threshold.
7. Optional: Click the **TEST** button to create a test alert, and to execute this alert action.

   For example, a test can be useful for checking whether email or SNMP is configured correctly.
8. Optional: To specify additional actions for this threshold event alert, click the add icon ⊕ in the **Alert actions** section title.

   Select the action and specify arguments.
9. Click **APPLY** at the top right of the dialog box.

Next Steps

- To modify a threshold alert, double-click the alert, or hover over the alert and click the edit icon ✎ and make the changes. Then click **APPLY**.
- To delete a threshold alert, hover over the alert, click the delete icon 🗑 and confirm that you want to delete this alert.

Related Topics

- Threshold Alert Properties
- Alert Action Types

Adding a Threshold Alert (CLI)

Use this procedure to specify an alert for a specific Analytics statistic threshold event.

Before You Begin

To perform this procedure, you must have the alert `allow_configure` authorization.

1. Go to `configuration alerts thresholds` and enter the `create` command.

   ```
   hostname:configuration alerts thresholds> create
   ```

2. Enter `get` to see the list of threshold alert properties.
3. Set threshold alert properties.
   - You must set at least `statname` and `limit`. Other properties that must be set have default values.
   - To see the list of statistics for which you can set threshold alerts, enter `set statname=` followed by a tab character.
4. Enter `commit`.
5. Copy the UUID in the commit message.
6. Create an alert action for this threshold alert.

   Follow the procedure Adding an Action for a Specific Threshold Alert (CLI), starting with step 2.

Next Steps

- To modify a threshold alert, enter `select threshold-###`, set the threshold alert properties, and enter `commit`.
- To specify additional actions for this event alert, follow the procedure Adding an Action for a Specific Threshold Alert (CLI).
- To delete a threshold alert, enter `destroy threshold-###` and enter `y` to confirm.

Related Topics

- Threshold Alert Properties
- Alert Action Types
- Adding an Action for a Specific Threshold Alert (CLI)

Threshold Alert Properties

Use the following properties to specify the statistic, define the threshold, and define when alert actions will be executed for threshold alerts.

Table 2-40 Threshold Alert Properties

BUI Property	CLI Property	Description
statistic name	`statname`	The statistic to monitor.

Table 2-40 (Cont.) Threshold Alert Properties

BUI Property	CLI Property	Description

exceeds
falls below
	`type=normal`
	`type=inverted`
		How to compare the threshold value to the current statistic value.

percent	`limit`	The integer percent or the number of bytes, operations, accesses, or requests per second.
for at least Select time unit.	`minpost` The time unit is seconds.	Integer length of time that the statistic value must remain in the threshold condition before the alert action is executed.
only between	`window_start` `window_end`	The window of time during which to execute this alert action. Select times from 00:00 through 23:30 UTC. To execute this alert action any time the conditions are met, specify `none` as either the start time or the end time.
only during	`days`	Which days to send these alerts. Choices are `all days`, `weekdays`, or `weekends`.
Repost alert every ... while this condition persists. Select time unit.	`frequency` The time unit is seconds.	Integer length of time between re-executing the alert action while the statistic value remains in the threshold condition.
Also post alert when this condition clears for at least ... Select time unit.	`minclear` The time unit is seconds.	Integer length of time that the statistic value must remain outside the threshold condition before a followup alert action is executed.

Configuring Certificates

This section describes the use of public key certificates. Public key certificates and their trust chains provide a mechanism to digitally identify a system without having to manually exchange any secret information.

A public key certificate is a blob of data that encodes a public key value, some information about the generation of the certificate, such as a name and who signed it, a hash or checksum of the certificate, and a digital signature of the hash. Together, these values form the certificate. The digital signature ensures that the certificate has not been modified. See Certificate Properties.

The appliance supports customer-owned certificates. The life cycle of a certificate starts with generating a certificate signing request (CSR). The CSR is then sent to the certificate authority (CA) for signature. After the signed certificate is returned from the CA, it can be installed on the appliance. If a certificate is signed by a non-root CA, you must also obtain certificates from the second- and higher-level CAs.

You can manage the following two types of certificates:

- System certificates identify the current system.
- Trusted certificates identify remote systems.

To manage system certificates, use the following tasks:

- Creating a New System Certificate - BUI, CLI
- Uploading CA Certificates from Non-root CAs - BUI, CLI

- Viewing CSR and System Certificate Details - BUI, CLI
- Destroying a CSR or System Certificate - BUI, CLI
- Setting the Appliance or Default System Certificate - BUI, CLI

To manage trusted certificates, use the following tasks:

- Uploading a Trusted Certificate - BUI, CLI
- Viewing Trusted Certificate Details - BUI, CLI
- Destroying a Trusted Certificate - BUI, CLI
- Assigning a Certificate to a Service - BUI, CLI

To use HTTP Strict Transport Security (HSTS) in conjunction with certificates, see the following topic: HTTP Strict Transport Security

Certificate Properties

This section lists the properties that describe system certificates, trusted certificates, and certificate signing requests (CSRs). These property values are read-only unless you are creating a CSR. These properties are optional and read-only unless specified. The property value format is a text string unless specified.

The following properties specify information about a certificate:

`comment`
Specifies an optional comment.

`dns`
Specifies a list of DNS names for which this certificate is issued. By default, this property value specifies the DNS names for this system's IP addresses. This value enables the client to verify that you have reached the intended system. This value is read-only except when creating a CSR.

`dirname`
Specifies a list of LDAP/X.500-style distinguished names for which the certificate is issued. This value is read-only except when creating a CSR.

`ip`
Specifies a list of IP addresses for which this certificate is issued. By default, this property specifies the system's IP addresses. This value enables the client to verify that you have reached the intended system. This value is read-only except when creating a CSR.

`notafter`
Specifies a time after which the certificate cannot be used. The timestamp is formatted as:

```
YYYY-[M]M-[D]D HH:mm[:ss]
```

Note that this value is always read-only and cannot be specified as part of a CSR.

`notbefore`
Specifies a time before which the certificate cannot be used. The timestamp is formatted as:

```
YYYY-[M]M-[D]D HH:mm[:ss]
```

Note that this value is always read-only and cannot be specified as part of a CSR.

serialnumber
Specifies the serial number of the certificate. Note that this value is always read-only and is present only in a certificate, not in a CSR.

sha1fingerprint
Specifies the SHA1 fingerprint of the certificate. This fingerprint value is automatically generated. The format is a list of hexadecimal pairs that are separated by colons. Note that this value is always read-only and is present only in a certificate, not in a CSR.

sha256fingerprint
Specifies the SHA256 fingerprint of the certificate. This fingerprint value is automatically generated. The format is a list of hexadecimal pairs that are separated by colons. Note that this value is always read-only and is present only in a certificate, not in a CSR.

type
Specifies the type of this entry. This value is read-only and automatically generated. Valid values are:

- `cert` specifies that the entry is a certificate
- `CA` specifies that the entry is a CA-certificate
- `request` specifies that the entry is a CSR
- `key` specifies that the entry is a key

uri
Specifies a list of universal resource identifiers (URIs) for which this certificate is issued. There is no default value. This value enables the client to verify that you have reached the intended system. This value is read-only except when creating a CSR.

uuid
Specifies the universally unique identifier (UUID) for this entry. This value is read-only and automatically generated.

The following read-only property values provide information about the certificate issuer and are under the control of the certificate authority (CA). You can use the following information to find the certificate.

- `issuer_commonname` - Specifies the certificate issuer's common name.
- `issuer_countryname` - Specifies the country name.
- `issuer_emailaddress` - Specifies the email address.
- `issuer_localityname` - Specifies the locality name, such as a city or town.
- `issuer_organizationalunitname` - Specifies the organizational unit name.
- `issuer_organizationname` - Specifies the organization name.
- `issuer_stateorprovincename` - Specifies the state or province name.

The following property values provide information about the certificate subject. You can use the following information to find the subject's certificate.

When you create a CSR to obtain a host certificate, you supply the following information about the host.

- `subject_commonname` - Specifies the certificate subject's common name. By convention, this value is the system's canonical DNS name. When you create a CSR, you must specify this property value for the host certificate.

- `subject_countryname` - Specifies the country name.
- `subject_emailaddress` - Specifies the email address.
- `subject_localityname` - Specifies the locality name, such as a city or town.
- `subject_organizationalunitname` - Specifies the organizational unit name.
- `subject_organizationname` - Specifies the organization name.
- `subject_stateorprovincename` - Specifies the state or province name.

The following properties control the creation of encryption keys for CSRs and for the certificates that are generated from them:

`key_type`
Specifies the encryption type. You must specify one of the following property values:

- `RSA` (for Rivest-Shamir-Adleman) is the default value
- `EC` (for Elliptic Curve)

`key_bits`
Specifies a key size. This value depends on the value of `key_type`.

- When `key_type=RSA`:
  - Reports a key size.
  - When creating a CSR, requests a key size, which is an even number from 2048 to 4096. The default value is 2048.
- When `key_type=EC`: Reports the key size for an EC certificate, CSR, or key.

`key_curve`
Specifies a list of EC curve values. Note that this list is subject to change over time. Requests or reports a particular curve. The default value is `prime256v1` (P-256). This value is used only when `key_type=EC`.

Creating a New Server Certificate (BUI)

Use this procedure to create a new server certificate.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **System** tab.
2. Create a new CSR.

   Either add a new CSR or copy an existing CSR.

   - To add a new CSR, click the add item icon ⊕ .
   - To create a new CSR based on an existing CSR or certificate, hover over an existing entry and click the copy icon ⊞ .
3. Complete the fields in the **New Certificate Request** dialog box.
4. Click **CREATE**.
5. Save the new CSR.

   When prompted to open the CSR or save it, save the CSR either now or later.

   - To save the CSR now, select **Save File** and click **OK**.
   - To save the CSR later, do the following:

   a. In the open or save dialog box, click **Cancel**.
   b. Hover over the entry in the table, and click the download icon ⤓.
6. Transfer the CSR to your CA in the prescribed manner.
7. Upload the signed certificate.

   After you receive the signed certificate from the CA, do the following:

   a. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **System** tab.
   b. Click the upload icon ⤒.
   c. Browse to the signed certificate and select it.
   d. Click **UPLOAD**.

Creating a New Server Certificate (CLI)

Use this procedure to create a new server certificate.

1. Go to `configuration settings certificates system`.
2. Create a new Certificate Signing Request (CSR).

   Either add a new CSR or copy an existing CSR.

   - To add a new CSR, enter the `create` command.
   - To create a new CSR based on an existing CSR or certificate, do the following:

     a. Enter the `list` command to view the certificates table.
     b. Enter the `clone` *`cert`* command, where *cert* is a value from the `CERT` column of the table.

     ```
     hostname:configuration settings certificates system> clone cert-001
     ```

3. Complete the CSR form.

   ```
   hostname:configuration settings certificates system (uncommitted)> get
               subject_commonname = hostname.us.example.com
         subject_organizationname = (unset)
   subject_organizationalunitname = (unset)
             subject_localityname = (unset)
      subject_stateorprovincename = (unset)
              subject_countryname = (unset)
             subject_emailaddress = (unset)
                              dns = hostname.us.example.com
                               ip = 192.0.2.174
                              uri = (unset)
                          dirname = (unset)
                          comment = (unset)
                         key_type = RSA
                         key_bits = 2048
                        key_curve = prime256v1 / P-256 (unused)
   hostname:configuration settings certificates system (uncommitted)> set
   comment="test"
                          comment = test (uncommitted)
   hostname:configuration settings certificates system (uncommitted)> commit
   ```

4. View the CSR.

a. Enter the `list` command to see your new CSR in the table.

```
hostname:configuration settings certificates system> list
CERT     TYPE SUBJECT COMMON NAME       ISSUER COMMON NAME        NOT AFTER
cert-002 req  hostname.us.example.com
cert-001 cert hostname.us.example.com   CA                        2023-1-25
cert-000 cert 3ebff8d2-58f6-4de4-a2c... 3ebff8d2-58f6-4de4-a2c... 2038-1-19
```

b. Enter the `dump cert` command, where `cert` is your new CSR in the table.

```
hostname:configuration settings certificates system> dump cert-002
-----BEGIN CERTIFICATE REQUEST-----
...
-----END CERTIFICATE REQUEST-----
```

5. Copy the CSR and transfer it to your CA in the prescribed manner.

6. Import the signed certificate.

   After you receive the signed certificate from the CA, do the following:

   a. Go to `configuration settings certificates system`.

   b. Enter the `import` command.

   c. At the prompt, paste the signed certificate.

```
hostname:configuration settings certificates system> import
("." to end)>
                  -----BEGIN CERTIFICATE-----
...
("." to end)>
                  -----END CERTIFICATE-----
("." to end)>
                  .
```

   The certificate replaces the CSR.

7. Verify the imported certificate.

   a. Enter the `list` command to see your new signed certificate in the table.

   b. Enter `select` *`cert`* and then enter `get` to view the properties of the certificate.

   c. Enter `done` and then enter `dump` *`cert`* to view the certificate.

Uploading CA Certificates from Non-root CAs (BUI)

If your server certificate is signed by a non-root CA, you need to obtain certificates for the second-level and higher-level CAs also. After obtaining these CA certificates, use this procedure to upload them.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **System** tab.

2. Click the upload icon ⤒ .

3. Browse to the signed certificate and select it.

4. Click **UPLOAD**.

5. Repeat steps 3 and 4 for each signed certificate.

Uploading CA Certificates from Non-root CAs (CLI)

If your server certificate is signed by a non-root CA, you also need to obtain certificates for the second-level and higher-level CAs. After obtaining these CA certificates, use this procedure to upload them.

1. Go to `configuration settings certificates system`.
2. Enter the `import` command.

   At the prompt, paste the signed certificate.

   ```
   hostname:configuration settings certificates system> import
   ("." to end)> -----BEGIN CERTIFICATE-----
     ...
   ("." to end)> -----END CERTIFICATE-----
   ("." to end)> .
   ```

3. Repeat step 2 for each signed certificate.
4. To check the imported certificates, use the `list` command to view the table of certificates.

   ```
   hostname:configuration settings certificates system> list
   CERT     TYPE SUBJECT COMMON NAME       ISSUER COMMON NAME        NOT AFTER
   cert-002 cert hostname.us.example.com   CA                        2023-1-25
   cert-001 cert 3ebff8d2-58f6-4de4-a2c... 3ebff8d2-58f6-4de4-a2c... 2038-1-19
   ```

Viewing CSR and Certificate Details (BUI)

Use this procedure to view CSR and certificate details.

A system certificate can be an automatically generated domain- or IP-address-based certificate, an automatically-generated ASN-based certificate, or a CA-signed certificate.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **System** tab.

   If you have not deleted them, you should see at least one automatically generated certificate based on the domain or IP address, and exactly one automatically-generated certificate based on the Appliance Serial Number (ASN) UUID.

2. Hover over an existing entry, and click its information icon ⓘ .

   The values of the **Subject Common Name**, **Issuer Common Name**, and **DirName** (distinguished name) are the ASN UUID. For a cluster, **DirName** includes the ASN UUID of each peer.

3. When finished, click **OK** to close the **Details** window.

Viewing CSR and Certificate Details (CLI)

Use this procedure to view CSR and certificate details.

A system certificate can be an automatically-generated domain- or IP-address-based certificate, an automatically-generated ASN-based certificate, or a CA-signed certificate.

1. Go to `configuration settings certificates system`.
2. Enter the `list` command.

If you have not deleted them, you should see at least one automatically generated certificate based on the domain or IP address, and exactly one automatically generated certificate based on the Appliance Serial Number (ASN) UUID.

```
hostname:configuration settings certificates system> list
CERT     TYPE SUBJECT COMMON NAME   ISSUER COMMON NAME     NOT AFTER
cert-002 cert alice.example.com...  alice.example.com...   2038-1-19
cert-001 cert 17f5fdce-6d64-4736... 17f5fdce-6d64-4736-... 2038-1-19
```

3. Use the `get` command to view the details of a CSR or certificate.

   - The following is an example of an automatically generated conventional certificate.

     ```
     hostname:configuration settings certificates system> select cert-002
     hostname:configuration settings certificates system cert-002> get
                   uuid = uuid
     subject_commonname = alice.example.com
      issuer_commonname = alice.example.com
                    dns = alice.example.com,alice,ip-addr
                     ip = 192.0.2.2
                    uri = https://alice.example.com:215,https://alice:215,https://ip-
     addr
                comment = Automatically generated
              notbefore = 2006-2-15 18:00
               notafter = 2038-1-19 03:14:07
           serialnumber = 59:8A:73:7B:00:00:00:27
        sha1fingerprint = 0A:14:26:ED:C7:43:0D:30:33:98:87:24:C5:9B:A2:52:55:FE:B1:D7
               key_type = RSA
               key_bits = 2048
     ```

   - The following is an example of a CSR.

     ```
                            uuid = uuid
              subject_commonname = alice.example.com
        subject_organizationname = Example Corp, Inc
            subject_localityname = Exampleton
     subject_stateorprovincename = CA
             subject_countryname = US
                             dns = alice.example.com
                              ip = 192.0.2.2
                        key_type = EC
                        key_bits = 256
                       key_curve = prime256v1 / P-256
     ```

   - The following is the CA-signed certificate that results from the preceding CSR.

     ```
                            uuid = uuid
              subject_commonname = alice.example.com
        subject_organizationname = Example Corp, Inc
            subject_localityname = Exampleton
     subject_stateorprovincename = CA
             subject_countryname = US
               issuer_commonname = Most Trusted Certificate
         issuer_organizationname = Totally Trustworthy Certificates, Inc
             issuer_localityname = Trustville
      issuer_stateorprovincename = AK
              issuer_countryname = US
                             dns = alice.example.com
                              ip = 192.0.2.2
                       notbefore = 2021-3-16 17:51:19
                        notafter = 2027-3-15 08:32:00
                    serialnumber = 4F
                 sha1fingerprint =
     62:FB:29:84:8C:3E:0E:C6:D2:49:88:38:F2:53:12:8D:A5:F9:96:88
     ```

```
 key_type = EC
 key_bits = 256
key_curve = prime256v1 / P-256
```

Destroying a CSR or Certificate (BUI)

Use this procedure to destroy a CSR or certificate.

> **Note:**
>
> Destroying a CSR also destroys the associated private key. Therefore, importing a certificate derived from that CSR will not be possible. Destroying a certificate also destroys the associated private key. Therefore, re-importing that certificate will not be possible.

After an ASN-based certificate is destroyed, a new ASN-based certificate is generated automatically without restart of the appliance software.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **System** tab.
2. Hover over an existing entry, and click the trash icon 🗑 .
3. Click **DESTROY**.

Destroying a CSR or Certificate (CLI)

Use this procedure to destroy a CSR or certificate.

> **Note:**
>
> Destroying a CSR also destroys the associated private key. Therefore, importing a certificate derived from that CSR will not be possible. Destroying a certificate also destroys the associated private key. Therefore, re-importing that certificate will not be possible.

After an ASN-based certificate is destroyed, a new ASN-based certificate is generated automatically without restart of the appliance software.

1. To view all certificate entries, go to the `configuration settings certificates system` context, and enter the `list` command.
2. To destroy a CSR or certificate, use the `destroy` command.

   ```
   hostname:configuration settings certificates system> destroy cert-002
   Caution: Destroying a certificate issued by a certificate authority also
   destroys the associated private key. Re-importing the certificate will not be
   possible.
   Destroy appliance certificate? (Y/N) y
   ```

Setting the Appliance Default Certificate (BUI)

Use this procedure to set the default certificate for this appliance.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **System** tab.
2. From the **System default certificate** menu, select the certificate that you want to set as the default for this appliance.
3. Click **APPLY**.

Setting the Appliance Default Certificate (CLI)

Use this procedure to set the default certificate for this appliance.

1. To view all certificate entries, go to the `configuration settings certificates system` context, and enter the `list` command.
2. To set a certificate as the default, use the `default` command.

```
hostname:configuration settings certificates system> set default=cert-002
                   default= cert-002 (uncommitted)
hostname:configuration settings certificates system> commit
```

Uploading Trusted Certificates (BUI)

Use this procedure to upload a trusted certificate.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **Trusted** tab.
2. Click the upload icon ⤒ .
3. Browse to the signed certificate and select it.
4. Click **UPLOAD**.
5. Repeat steps 3 and 4 for each signed certificate.
6. Check the imported certificates.

   See Viewing Trusted Certificate Details (BUI).

Uploading Trusted Certificates (CLI)

Use this procedure to upload a trusted certificate.

1. Go to `configuration settings certificates trusted`.
2. Enter the `import` command.

   At the prompt, paste the signed certificate.

```
hostname:configuration settings certificates system> import
("." to end)> -----BEGIN CERTIFICATE-----
  ...
("." to end)> -----END CERTIFICATE-----
("." to end)> .
```

3. Repeat step 2 for each signed certificate.
4. Check the imported certificates.

   See Viewing Trusted Certificate Details (CLI).

Viewing Trusted Certificate Details (BUI)

Use this procedure to view trusted certificate details.

A trusted certificate can be an automatically generated domain- or IP-address-based certificate, an automatically-generated ASN-based certificate, or a CA-signed certificate.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **Trusted** tab.

   If you have not deleted them, you should see at least one automatically-generated certificate based on the domain or IP address, and exactly one automatically-generated certificate based on the Appliance Serial Number (ASN) UUID.

2. Hover over an existing entry, and click its information icon ⓘ .

   The values of the **Subject Common Name**, **Issuer Common Name**, and **DirName** (distinguished name) are the ASN UUID. For a cluster, **DirName** includes the ASN UUID of each peer.

3. When finished, click **OK** to close the **Details** window.

Viewing Trusted Certificate Details (CLI)

Use this procedure to view trusted certificate details.

A trusted certificate can be an automatically generated domain- or IP-address-based certificate, an automatically-generated ASN-based certificate, or a CA-signed certificate.

1. Go to `configuration settings certificates trusted`.
2. Use the `list` command to show the certificate list.

   If you have not deleted them, you should see at least one automatically-generated certificate based on the domain or IP address, and exactly one automatically-generated certificate based on the Appliance Serial Number (ASN) UUID.

   ```
   hostname:configuration settings certificates trusted> list
   CERT     TYPE SUBJECT COMMON NAME    ISSUER COMMON NAME     NOT AFTER
   cert-002 cert cd9dfcfe-b0d3-600f-... cd9dfcfe-b0d3-600f-... 2038-1-19
   cert-001 cert 3638ef45-ae8c-ec31-... 3638ef45-ae8c-ec31-... 2038-1-19
   ```

3. To view the details of a certificate, `select` the certificate, and use the `get` command.

Destroying a Trusted Certificate (BUI)

Use this procedure to destroy a trusted certificate. After an ASN-based certificate is destroyed, a new ASN-based certificate is generated automatically without restart of the appliance software.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **Trusted** tab.
2. Hover over an existing entry, and click the trash icon 🗑 .
3. Click **DESTROY**.

Destroying a Trusted Certificate (CLI)

Use this procedure to destroy a trusted certificate.

After an ASN-based certificate is destroyed, a new ASN-based certificate is generated automatically without restart of the appliance software.

1. To view all certificate entries, go to the `configuration settings certificates trusted` context, and enter the `list` command.
2. To destroy a certificate, use the following commands.

```
hostname:configuration settings certificates system> destroy cert-001
Caution: Destroying a certificate issued by a certificate authority also
destroys the associated private key. Re-importing the certificate will not be
possible.
Destroy appliance certificate? (Y/N) y
```

Assigning a Certificate to a Service (BUI)

Use this procedure to assign a certificate to a service, such as LDAP.

1. From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **Trusted** tab.
2. From the drop-down menu of system certificates, select the certificate that you want to assign.
3. Click the edit icon ✎ .
4. Select the service, such as **ldap**, from the list of services at the bottom of the page.

Assigning a Certificate to a Service (CLI)

Use this procedure to assign a certificate to a service, such as LDAP.

1. To view all certificate entries, go to the `configuration settings certificates trusted` context, and enter the `list` command.
2. Select the certificate to which you want to assign the service, and set the service's name.

```
hostname:configuration settings certificates trusted> select cert-001
hostname:configuration settings certificates trusted cert-001> set services=ldap
                    services= ldap (uncommitted)
hostname:configuration settings certificates trusted cert-001> commit
```

HTTP Strict Transport Security

HTTP Strict Transport Security (HSTS) allows only secure HTTPS connections, and not HTTP connections, for a specified period of time. Before using HSTS, familiarize yourself with HSTS prerequisites, understand browser behavior with HSTS enabled, and install a certificate signed by a certificate authority.

> **Note:**
>
> Failure to keep the certificate valid and appropriate could negate HSTS security advantages or could cause a browser to not connect with the appliance.

To enable HSTS, use the following procedure: Enabling HTTP Strict Transport Security - BUI, CLI

Enabling HTTP Strict Transport Security (BUI)

Use this procedure to enable HSTS for this appliance.

1. From the **Configuration** menu, select **Settings**, then **Certificates**.
2. Under **Security Settings**, select the check box to enable **HSTS**.
3. Optional: Set the maximum length of time that HSTS will remain enabled.

   Update the **HSTS Max Age**, including units.
4. Click **APPLY**.

Enabling HTTP Strict Transport Security (CLI)

Use this procedure to enable HSTS for this appliance.

1. Go to `configuration settings certificates security`.
2. Enter the `get` command to view security properties.

   ```
   hostname:configuration settings certificates security> get
                  hsts_enable = false
                 hsts_max_age = 730 days
   ```
3. Set the `hsts_enable` property to `true`.

   ```
   hostname:configuration settings certificates security> set hsts_enable=true
                  hsts_enable = true (uncommitted)
   ```
4. Optional: Set the maximum length of time that HSTS will remain enabled.

   Update the value of `hsts_max_age`, including units. Units can be `seconds`, `minutes`, `hours`, `days`, `weeks`, or `months`.
5. Enter `commit` to commit the changes.

Configuring SSL/TLS Versions and Ciphers

This section describes how to configure SSL/TLS protocol versions and ciphers that Oracle ZFS Storage Appliance uses to communicate with peer systems.

> **Caution:**
>
> To avoid service unavailability, keep the default settings unless otherwise needed or as instructed by Oracle Support.

SSL/TLS ciphers are algorithms used to encrypt and decrypt data as it is transmitted across the network. Configure the SSL/TLS versions and ciphers according to your site's security requirements. Ensure that the source and target systems are configured to support overlapping values.

As of software release OS8.8.67, TLSv1.0 and TLSv1.1 are not supported.

Starting with the OS8.8.69 release, TLS versions 1.2 and 1.3 are supported. The TLS 1.3 ciphers are separate from TLS 1.2 ciphers. If you enable TLS 1.2, you must enable at least one TLS 1.2 cipher. If you enable TLS 1.3, you must enable at least one TLS 1.3 cipher.

Future Oracle ZFS Storage Appliance releases might drop support for older versions and ciphers and versions. Systems that run older firmware might not support newer versions and ciphers. For two systems to communicate, they must share at least one version and at least one cipher. In particular, if one system supports only TLSv1.2, all systems must be configured to allow TLSv1.2.

To configure SSL/TLS, use the following tasks:

- Configuring SSL/TLS (BUI)
- Configuring SSL/TLS (CLI)

Configuring SSL/TLS (BUI)

To configure SSL/TLS versions and ciphers, use the following steps. The versions and at least one of the ciphers must be identical on all appliances that communicate with each other.

1. From the **Configuration** menu, select **Settings**, then **Peer**.
2. Click **Edit** next to **SSL/TLS versions and ciphers**.
3. Set the versions and ciphers, and click **OK**.

   The list of ciphers varies per the versions selected.

Configuring SSL/TLS (CLI)

To configure SSL/TLS versions and ciphers, use the following steps. The versions and at least one of the ciphers must be identical on all appliances that communicate with each other.

1. Go to `configuration settings peer` and enter `ls` to list the SSL/TLS versions and ciphers.

   The list of ciphers varies per the versions selected.
2. Enter the SSL/TLS versions using command `set tls_version` and the version name.

   ```
   hostname:configuration settings peer> set tls_version=TLSv1.2
                  tls_version = TLSv1.2 (uncommitted)
   ```
3. Enter the ciphers using command `set ciphers` and the cipher names, separated by commas.

   ```
   hostname:configuration settings peer> set
   ciphers=AES128-GCM-SHA256,ECDH-ECDSA-AES128-GCM-SHA256
                  ciphers =
   AES128-GCM-SHA256,ECDH-ECDSA-AES128-GCM-SHA256 (uncommitted)
   ```
4. Enter `commit`. To view the versions and ciphers, enter `show`.

   ```
   hostname:configuration settings peer> commit
   hostname:configuration settings peer> show
   Properties:
                  tls_version = TLSv1.2
                      ciphers =
   AES128-GCM-SHA256:ECDH-ECDSA-AES128-GCM-SHA256
   hostname:configuration settings peer>
   ```

Configuring Password Complexity

An administrative user with sufficient privileges can set password complexity rules for all local users. Subsequently, when local users set or change their passwords, they are constrained by these complexity rules; existing passwords are not affected.

To configure password complexity for local users, see the following sections:

- Setting Password Complexity - BUI, CLI
- Password Complexity Properties

For information about user configuration, including roles, authorizations, and exceptions, see Configuring Users.

Automatic Account Locking

The automatic account locking feature, by default, locks a local user's account for five minutes after five failed password entries. This feature is available for local administrative or data users, but not for directory-based accounts or directory roles. Also, this feature cannot be used to manually lock or unlock an account. However, a locked account can be unlocked before the automatic unlock time expires by resetting the password.

Configure or disable this feature when setting password complexity properties.

Setting Password Complexity (BUI)

Use the following procedure to set password complexity for local users.

Before You Begin

For the administrative user's role or exceptions, if all services are not authorized for the services scope, add the `configure` authorization for the password service.

1. From the **Configuration** menu, select **Settings**, then **Passwords**.
2. Modify the properties with values described in Password Complexity Properties.
3. Click **APPLY**.

Related Topics

- Password Complexity Properties
- Configuring Users

Setting Password Complexity (CLI)

Use the following procedure to set password complexity for local users.

Before You Begin

For the administrative user's role or exceptions, if all services are not authorized for the services scope, add the `configure` authorization for the password service.

1. Go to `configuration settings passwords`.
2. Enter `show` to view the properties and their values.
3. Set the properties with values described in Password Complexity Properties.

4. Enter `commit`.

Related Topics

- Password Complexity Properties
- Configuring Users

Password Complexity Properties

The following table describes the properties for configuring password complexity for local users.

Table 2-41 Password Complexity Properties

BUI Name	CLI Name	Description
Minimum length	`passlength`	Minimum character length of the password. Must be a positive integer. Default value: `8`.
Minimum number of letters	`min_letters`	Minimum number of alphabetic characters. Default value: `0`.
Minimum number of upper case letters	`min_upper`	Minimum number of upper case alphabetic characters. Default value: `0`.
Minimum number of lower case letters	`min_lower`	Minimum number of lower case alphabetic characters. Default value: `0`.
Minimum number of digits	`min_digit`	Minimum number of numeric characters. Default value: `0`.
Minimum number of punctuation characters	`min_punctuation`	Minimum number of punctuation characters. Note that some versions of Oracle ILOM do not support colons and spaces in the root password. Default value: `0`.
Maximum number of consecutive repeating characters	`max_repeats`	Maximum number of consecutive repeating characters. Default value: `0`.
Password must not be a circular shift of username	`namecheck`	Controls whether the password can be based on the username. Default value: `true`.

The following table describes the properties for configuring automatic account locking for local users.

Table 2-42 Automatic Account Locking Properties

BUI Name	CLI Name	Description
Lock after failed login attempts	`lock_after_retries`	Indicates if the account will be locked after the set number of failed password entry attempts. Default value: `true`.
Login attempts	`retries`	Number of failed password entry attempts before the account is locked. Default value: `5`.

Table 2-42 (Cont.) Automatic Account Locking Properties

BUI Name	CLI Name	Description
Unlock after	`auto_unlock_time`	Number of seconds/minutes/hours/days that the account remains locked after meeting the retries threshold. Default value: `5 minutes`.

3 Appliance Services

Appliance services are easily managed from the BUI **Configuration: Services** screen, or the CLI `configuration services` context.

Use the following tasks for viewing and managing appliance services:

- Viewing a Service in the BUI
- Selecting a Service in the CLI
- Enabling a Service - BUI, CLI
- Disabling a Service - BUI, CLI
- Viewing Service States in the CLI
- Viewing Service Help in the CLI
- Setting Service Properties - BUI, CLI
- Viewing Service Logs - BUI, CLI
- List of Available Appliance Services
- Required Service Ports

For information about configuring an individual service, select the service from the following table.

Data Services	Directory Services	System Settings	Remote Access
NFS	NIS	DNS	SSH
iSCSI	LDAP	IPMP	RESTful API
SMB	Active Directory	NTP	HTTPS
FTP	Identity Mapping	Phone Home	
HTTP	RADIUS	Dynamic Routing	
NDMP		Service Tags	
Remote Replication		SMTP	
Shadow Migration		SNMP	
SFTP		Syslog	
SRP		System Identity	
TFTP		Kerberos	
Virus Scan		Retention	
Cloud			

Managing Services

For information about managing appliance services, use the following tasks:

- Viewing a Service in the BUI

- Selecting a Service in the CLI
- Enabling a Service - BUI, CLI
- Disabling a Service - BUI, CLI
- Viewing Service States in the CLI
- Viewing Service Help in the CLI
- Setting Service Properties - BUI, CLI
- Viewing Service Logs - BUI, CLI
- List of Available Appliance Services
- Required Service Ports

Viewing a Service in the BUI

Use the following procedure to view a service.

1. From the **Configuration** menu, select **Services**.
2. To view or edit the properties for a specific service, hover over the service status icon that is to the left of the service name.

   The status icon turns into an arrow icon ➔ .
3. Click the arrow icon ➔ to display the properties screen for the selected service.
4. In any of the services screens, you can show a side panel of all services by clicking the small arrow icon to the left of the **Services** title (near the top left of each screen). Click this icon again to hide the list.

Related Topics

- Enabling a Service (BUI)
- Setting Service Properties (BUI)
- List of Available Appliance Services
- Required Service Ports

Selecting a Service in the CLI

After you select a service, you can view its state, enable it, disable it, and set its properties.

1. Go to `configuration services`.
2. Select a service by entering its name. For example, enter `nis`:

```
hostname:configuration services> nis
hostname:configuration services nis>
```

Related Topics

- Enabling a Service (CLI)
- Setting Service Properties (CLI)
- List of Available Appliance Services
- Required Service Ports

Enabling a Service (BUI)

Use the following procedure to enable a service that is not online.

1. From the **Configuration** menu, select **Services**.
2. Click the power icon ⏻ to bring the service online .

Related Topics

- Disabling a Service (BUI)
- Setting Service Properties (BUI)
- List of Available Appliance Services
- Required Service Ports

Enabling a Service (CLI)

Use the following procedure to enable a service that is not online.

1. Go to `configuration services`.
2. Select a service, then enter the `enable` command to enable the service.

```
hostname:configuration services> nis
hostname:configuration services nis> enable
```

Related Topics

- Disabling a Service (CLI)
- Setting Service Properties (CLI)
- List of Available Appliance Services
- Required Service Ports

Disabling a Service (BUI)

Use the following procedure to disable a service that is online.

1. From the **Configuration** menu, select **Services**.
2. Click the power icon ⏻ to take the service offline .

Related Topics

- Enabling a Service (BUI)
- List of Available Appliance Services
- Required Service Ports

Disabling a Service (CLI)

Use the following procedure to disable a service that is online.

1. Go to `configuration services`.
2. Select the service, then enter the `disable` command to disable it.

```
hostname:configuration services> nis
hostname:configuration services nis> disable
```

Related Topics

- Enabling a Service (CLI)
- List of Available Appliance Services
- Required Service Ports

Viewing Service States in the CLI

Use the following procedure to view service states.

1. Go to `configuration services`.
2. Enter the `show` command to list the current state of all services.
3. To view the state of an individual service, select the service and then enter `show`.

```
hostname:configuration services> nis
hostname:configuration services nis> show
Properties:
                      <status> = online
                        domain = example
                     broadcast = true
                     ypservers =
```

Related Topics

- Enabling a Service (CLI)
- Setting Service Properties (CLI)
- List of Available Appliance Services
- Required Service Ports

Viewing Service Help in the CLI

Use the following procedure to display available commands for a service.

1. Go to `configuration services`.
2. Select the service and enter `help`.

```
hostname:configuration services> nis
hostname:configuration services nis> help
Subcommands that are valid in this context:

   help [topic]         => Get context-sensitive help. If [topic] is specified,
                           it must be one of "builtins", "commands", "general",
                           "help", "script" or "properties".

   show                 => Show information pertinent to the current context

   commit               => Commit current state, including any changes

   done                 => Finish operating on "nis"

   enable               => Enable the nis service

   disable              => Disable the nis service
```

```
get [prop]          => Get value for property [prop]. ("help properties"
                       for valid properties.) If [prop] is not specified,
                       returns values for all properties.

set [prop]          => Set property [prop] to [value]. ("help properties"
                       for valid properties.) For properties taking list
                       values, [value] should be a comma-separated list of
                       values.
```

Related Topics

- Setting Service Properties (CLI)
- List of Available Appliance Services
- Required Service Ports

Setting Service Properties (BUI)

The **Configuration: Services** screens allow you to view and modify the services. The following table describes the icons and buttons in the services screens.

Icon	Description
➔	Go to the service screen to configure properties and view logs. This button appears when you hover over a service.
▟	The service is enabled and working normally.
●	The service is offline or disabled.
●	The service has a problem and requires operator attention.
⏻	Enables or disables the service.
⟳	Restarts the service.
⏻	Enable/disable not available for this service.
⟳	Restarts the currently unavailable service. You must enable the service first.

1. From the **Configuration** menu, select **Services**.
2. Double-click on a service name.
3. Change the properties, and then click **APPLY**.

   To reset properties, click **REVERT**.

Related Topics

- Enabling a Service (BUI)
- List of Available Appliance Services
- Required Service Ports

Setting Service Properties (CLI)

Use the following procedure to define properties for a service. Property names are similar to their names in the BUI, but CLI names are usually shorter and sometimes abbreviated.

1. Go to `configuration services`.
2. Select a service and enter `show` to view the list of properties you can set for that service, along with their current values.

   ```
   hostname:configuration services> nis
   hostname:configuration services nis> show
   Properties:
                         <status> = online
                           domain =
                        broadcast = true
                        ypservers =
   ```

3. Use the `set` command to set the properties.

   ```
   hostname:configuration services nis> set domain="mydomain"
                           domain = mydomain (uncommitted)
   ```

4. After setting the properties, enter `commit` to save and activate the new configuration.

   ```
   hostname:configuration services nis> commit
   hostname:configuration services nis> show
   Properties:
                         <status> = online
                           domain = mydomain
                        broadcast = true
                        ypservers =
   ```

Related Topics

- Enabling a Service (CLI)
- List of Available Appliance Services
- Required Service Ports

Viewing Service Logs (BUI)

Some services provide service logs with information to help you diagnose service issues. If a **Logs** button exists in the top right of a service screen, that service provides a log. Service logs can provide the times when a service changed state, and the error messages from the service. Log content is specific to each individual service and is subject to change.

1. From the **Configuration** menu, select **Services**, and double-click on a service name.
2. Click the **Logs** button at the top right of a service screen.

   The following are common example messages:

Example Log Message	Description
`Executing start method`	The service is starting up.
`Method "start" exited with status 0`	The service reported a successful start (`0` == success).
`Method "refresh" exited with status 0`	The service successfully refreshed its configuration based on its service settings.
`Executing stop method`	The service is being shut down.
`Enabled`	The service state was checked to see if it should be started (such as during system boot), and it was found to be in the enabled state.

Example Log Message	Description
`Disabled`	The service state was checked to see if it should be started (such as during system boot), and it was found to be in the disabled state.

The following log example is from the NTP service:

```
[ Oct 15 21:05:31 Enabled. ]
[ Oct 15 21:07:37 Executing start method (...). ]
[ Oct 15 21:13:38 Method "start" exited with status 0. ]
```

The first log event in the example shows that the system was booted at 21:05. The second entry at 21:07:37 records that the service began startup, which completed at 21:13:38. Due to the nature of NTP and system clock adjustment, this service can take minutes to complete startup, as shown by the log.

Related Topics

- List of Available Appliance Services
- Required Service Ports

Viewing Service Logs (CLI)

You cannot view service logs from the CLI. Use the BUI as described in Viewing Service Logs (BUI).

List of Available Appliance Services

This section lists the available appliance services, along with short descriptions and port information. Certain services are always on and cannot be disabled, as described in the following table.

Table 3-1 Data Services

Service	Description	Ports Used
NFS	Filesystem access via the NFSv3, NFSv4.0, and NFS v4.1 protocols	111 and 2049
iSCSI	LUN access via the iSCSI protocol	3260 and 3205
SMB	Filesystem access via the SMB protocol	SMB-over-NetBIOS 139 SMB-over-TCP 445 NetBIOS Datagram 138 NetBIOS Name Service 137
FTP	Filesystem access via the FTP protocol	21
HTTP	Filesystem access via the HTTP protocol	80
NDMP	NDMP host service	10000
Remote Replication	Remote replication	216 and 217
Shadow Migration	Shadow data migration	
SFTP	Filesystem access via the SFTP protocol	218

Table 3-1 (Cont.) Data Services

Service	Description	Ports Used
SRP	Block access via the SRP protocol	
TFTP	Filesystem access via the TFTP protocol	
Virus Scan	Filesystem virus scanning	
Cloud	Cloud data service for cloud backups	

> **Note:**
>
> UIDs and GIDs from 0 to 99 are reserved by the operating system vendor for use in future applications. Their use by end-system users or vendors of layered products is not supported and may cause security-related issues with future applications.

Table 3-2 Directory Services

Service	Description	Ports Used
NIS	Authenticate users and groups from an NIS service	
LDAP	Authenticate users and groups from an LDAP directory	389
Active Directory	Authenticate users with a Microsoft Active Directory Server **Note:** Once enabled, this becomes an always-on service, and cannot be disabled.	
Identity Mapping	Map between Windows entities and UNIX IDs **Note:** Always-on service, cannot be disabled.	
RADIUS	Authenticate users with a RADIUS server **Note:** Once enabled, *all* directory users use this service.	

Table 3-3 Service Settings

Service	Description	Ports Used
DNS	Domain name service client **Note:** Always-on service, cannot be disabled.	53
Dynamic Routing	RIP and RIPng dynamic routing protocols	
IPMP	IP network multipathing for IP fail-over **Note:** Always-on service, cannot be disabled.	
Kerberos	Kerberos authentication	88
NTP	Network time protocol client	
Phone Home	Product registration and support configuration	8000 (It depends on the port opened up in the proxy)
Retention	File retention	
Service Tags	Product inventory support	6481

Table 3-3 (Cont.) Service Settings

Service	Description	Ports Used
SMTP	Configure outgoing mail server **Note:** Always-on service, cannot be disabled.	
SNMP	SNMP for sending traps on alerts and serving appliance status information	
Syslog	Syslog Relay for sending syslog messages on alerts and forwarding service syslog messages	
System Identity	System name and location **Note:** Always-on service, cannot be disabled.	

Table 3-4 Remote Access Services

Service	Description	Ports Used
SSH	SSH for CLI access	22
REST	RESTful API **Note:** Always-on service, cannot be disabled.	

Required Service Ports

To provide security on a network, you can deploy firewalls within your network architecture. Port numbers are used for creating firewall rules and to uniquely identify a transaction over a network by specifying the host and the service.

The following list shows the minimum ports required for creating firewall rules that allow full functionality of Oracle ZFS Storage Appliance:

Inbound Ports

- icmp/0-65535 (PING)
- tcp/1920 (EM)
- tcp/215 (BUI)
- tcp/22 (SSH)
- udp/161 (SNMP)

Outbound Ports

- tcp/80 (WEB)
- tcp/443 (SSL WEB)

Note:

An outbound port of tcp/443 is used for sending Phone Home messages, uploading support bundles, and update notifications. For replication, use Generic Routing Encapsulation (GRE) tunnels when possible. This lets traffic run on the back end interfaces and avoid the firewall where traffic could be slowed. If GRE tunnels are not available on the NFS core, you must run replication over the front end interface. In this case, port 216 and port 217 must also be open.

Configuring Services

For information about configuring a service, select one of the services from the following table.

Data Services	Directory Services	System Settings	Remote Access
NFS	NIS	DNS	SSH
iSCSI	LDAP	IPMP	RESTful API
SMB	Active Directory	NTP	HTTPS
FTP	Identity Mapping	Phone Home	
HTTP	RADIUS	Dynamic Routing	
NDMP		Service Tags	
Remote Replication		SMTP	
Shadow Migration		SNMP	
SFTP		Syslog	
SRP		System Identity	
TFTP		Kerberos	
Virus Scan		Retention	
Cloud			

Related Topics

List of Available Appliance Services

Active Directory Configuration

The Active Directory (AD) service provides access to a Microsoft Active Directory database, which stores information about users, groups, shares, and other shared objects.

The AD service has two modes: domain mode and workgroup mode. These modes dictate how SMB users are authenticated. When operating in domain mode, SMB clients are authenticated through the AD domain controller. In workgroup mode, SMB clients are authenticated locally as local users. See Configuring Users for more information about local users.

You can add an AD user as an appliance administrator. For more information, see Configuring Users.

Use one of the following methods to integrate with AD:

- Point your LDAP configuration at AD.

AD must contain UNIX data and the AD service must use UNIX-style names.

- Join an AD domain.

  Use an AD *name@domain* username. You do not need to configure LDAP or add UNIX data to AD.

To configure Active Directory, see the following sections:

- Joining an AD Domain (BUI)
- Joining a Workgroup (BUI)
- Automatically Configuring LDAP for an AD Domain (BUI)
- Configuring Active Directory (CLI)
- Active Directory Join Domain
- Active Directory Domains and Workgroups
- Active Directory Windows Support

Joining an AD Domain (BUI)

Use the following procedure to join an AD domain.

1. Optional: Configure an Active Directory site in the SMB context.
2. Optional: Configure a preferred domain controller in the SMB context.
3. Enable NTP, or ensure that the clocks of Oracle ZFS Storage Appliance and the domain controller are synchronized to within five minutes.
4. Ensure that your DNS infrastructure correctly delegates to the Active Directory domain, or set your domain controller's IP address as a DNS name server.

   You can specify up to three of your redundant domain controllers as DNS name servers.
5. From the **Configuration** menu, select **Services**, then **Active Directory**, and click **Join Domain**.
6. Configure the Active Directory domain, administrative user name, and administrative password.
7. Click **APPLY** to commit the configuration.

Related Topics

- Joining a Workgroup (BUI)
- Configuring Active Directory (CLI)
- Active Directory Join Domain
- Active Directory Domains and Workgroups
- Active Directory Windows Support

Joining a Workgroup (BUI)

Use the following procedure to join a workgroup.

1. From the **Configuration** menu, select **Services**, then **Active Directory**, and click **Join Workgroup**.
2. Enter the Windows workgroup name.
3. Select or clear the option to unconfigure the LDAP service, if the option is available.

- When LDAP has been configured with AD and the appliance is leaving the current AD domain, select this option to leave the AD domain permanently and to clear the LDAP configuration. After the LDAP service is unconfigured, it is automatically disabled.
- Clear this option if LDAP has been configured with AD and you want to later rejoin the current AD domain and retain the AD LDAP configuration.
- This option is unavailable when:
  - The LDAP service is not configured or it is configured for UNIX LDAP.
  - The LDAP service is configured for a different AD domain, not the current domain.
  - The appliance is not changing from domain mode to workgroup mode.

4. Click **APPLY** to commit the configuration.

Related Topics

- Joining an AD Domain (BUI)
- Configuring Active Directory (CLI)
- Active Directory Join Domain
- Active Directory Domains and Workgroups
- Active Directory Windows Support

Automatically Configuring LDAP for an AD Domain (BUI)

Use the following procedure to automatically configure the LDAP service for the currently joined Active Directory (AD) domain.

1. From the **Configuration** menu, select **Services**, and then **Active Directory**.
2. Ensure the appliance is joined to the appropriate AD domain or click **Join Domain** and join the domain.
3. Click **Configure LDAP for Use with This AD Domain**.
4. When the **SUCCESS** dialog box is displayed, click **OK** to accept the LDAP configuration or click **CUSTOMIZE** to view or change the configuration. For LDAP service configuration information, see LDAP Configuration.

   If only one LDAP server is configured, add more LDAP servers by clicking **CUSTOMIZE** and manually adding more servers. If the DNS service cannot discover any LDAP servers, click **CUSTOMIZE** and manually set up the servers in the LDAP service.

Related Topics

- Joining an AD Domain (BUI)
- Configuring Active Directory (CLI)
- Active Directory Join Domain
- Active Directory Domains and Workgroups
- Active Directory Windows Support
- LDAP Configuration

Configuring Active Directory (CLI)

Use the following procedure to configure Active Directory (AD).

1. Go to `configuration services ad`.

```
hostname:> configuration services ad
```

2. To view an existing configuration, enter `show`.

   ```
   hostname:configuration services ad> show
   Properties:
                        <status> = online
                            mode = domain
                          domain = eng.test.com
                          server = server-name.example.com
                     diagnostics = (unset)

   Children:
                             domain => Join an Active Directory domain
                          workgroup => Join a Windows workgroup
   ```

   Observe that the appliance is currently operating in the domain `eng.test.com`.

3. To join a new domain after the properties are configured, enter the following commands.

   When joining an AD domain, you must set the user and password each time you commit the node.

   ```
   hostname:> configuration services ad
   hostname:configuration services ad> domain
   hostname:configuration services ad domain> set domain=example.com
   hostname:configuration services ad domain> set user=Administrator
   hostname:configuration services ad domain> set password=(set)
   hostname:configuration services ad domain> commit
   hostname:configuration services ad domain> done
   hostname:configuration services ad> show
   Properties:
                        <status> = online
                            mode = domain
                          domain = example.com
                          server = server-name.example.com
                     diagnostics = (unset)
   ```

4. To configure the site and preferred domain controller in preparation for joining another domain, enter the following commands:

   ```
   hostname:configuration services ad> done
   hostname:> configuration services smb
   hostname:configuration services smb> set ads_site=sf
   hostname:configuration services smb> set pdc=192.0.2.21
   hostname:configuration services smb> commit
   hostname:configuration services smb> show
   Properties:
                        <status> = online
                    lmauth_level = 4
                             pdc = 192.168.3.21
                        ads_site = sf
   hostname:configuration services smb> done
   ```

5. To automatically configure the LDAP service for the currently joined AD domain, perform the following steps:

   a. Ensure the appliance is joined to the appropriate AD domain.

   b. Go to `configuration services ad` and enter `ldap`.

      ```
      hostname:> configuration services ad
      hostname:configuration services ad> ldap
      ```

c. Confirm your action to automatically configure LDAP by entering **y**.

```
Are you sure you want the system to automatically configure LDAP for this AD
domain? y
LDAP has been set up for use with this AD domain;
Please navigate to 'configuration services ldap' to view or customize.
```

d. To view the LDAP configuration, go to `configuration services ldap` and enter command `show`.

```
hostname:configuration services ad> cd ..
hostname:configuration services> ldap
hostname:configuration services ldap> show
```

e. To customize the LDAP service configuration, go to `configuration services ldap` and see LDAP Configuration.

```
hostname:configuration services ad> cd ..
hostname:configuration services> ldap
hostname:configuration services ldap>
```

6. To leave the domain mode, join a Windows workgroup, and to clear the LDAP configuration that has been set for the current domain, enter the following commands:

```
hostname:configuration services ad> workgroup
hostname:configuration services ad workgroup> set workgroup=WORKGROUP
hostname:configuration services ad workgroup> set unconfig_ldap=true
hostname:configuration services ad workgroup> commit
hostname:configuration services ad> show
Properties:
                    workgroup = WORKGROUP
                unconfig_ldap = true
hostname:configuration services ad workgroup> done
hostname:configuration services ad> show
Properties:
                     <status> = disabled
                         mode = workgroup
                    workgroup = WORKGROUP
```

Rules for setting property `unconfig_ldap`:

- When LDAP has been configured with AD and the appliance is leaving the current AD domain, set this property to true to leave the AD domain permanently and to clear the LDAP configuration. After the LDAP service is unconfigured, it is automatically disabled.
- Set this property to false if LDAP has been configured with AD and you want to later rejoin the current AD domain and retain the AD LDAP configuration.
- This property cannot be changed from false when the LDAP service is not configured. Additionally, this property cannot be changed from false, and the LDAP configuration remains intact, when:
  - The appliance is not changing from domain mode to workgroup mode.
  - LDAP is configured for a different AD domain, not the current domain.
  - LDAP is configured for UNIX LDAP.

Related Topics

- Joining an AD Domain (BUI)
- Joining a Workgroup (BUI)

- Automatically Configuring LDAP for an AD Domain (BUI)
- LDAP Configuration
- Active Directory Join Domain
- Active Directory Domains and Workgroups
- Active Directory Windows Support

Active Directory Join Domain

If an account does not already exist in Active Directory by default, a machine trust account for the system is automatically created in the default container for computer accounts (`cn`=Computers) as part of the domain join operation. The following users are allowed to perform domain join:

- **Domain administrator** - Can join any number of systems to the domain with machine trust accounts placed in any containers.
- **Delegated administrator with authority over one or more Organizational Units** - Can join any number of systems to a domain with machine account location designated in the Organizational Units for which they are responsible.
- **Normal user with machine accounts pre-staged by administrator** - Can join a system to the domain as pre-authorized by an administrator.
- **Normal user** - Normally authorized to join a limited number of systems.

The following properties for joining an Active Directory domain are available:

- **Active Directory Domain** - The fully-qualified name or NetBIOS name of an Active Directory domain
- **User** - An AD user who has credentials to create a computer account in Active Directory
- **Password** - The administrative user's password
- **Organizational Unit** - Specifies an alternative organizational unit in which the system's machine trust account will be created. The organizational unit is specified as a comma-separated list of one or more name-value pairs using the domain-relative distinguished name (DN) format, for example, `ou=innerOU,ou=outerOU`.
- **Use Pre-created Account** - If the system's account exists and the specified Organizational Unit is not the one that the account is in, use the pre-created account.

Related Topics

- Joining an AD Domain (BUI)
- Joining a Workgroup (BUI)
- Configuring Active Directory (CLI)
- Active Directory Domains and Workgroups
- Active Directory Windows Support

Active Directory Domains and Workgroups

The configurable property for joining a workgroup is Windows Workgroup.

Instead of enabling and disabling the service directly, the service is modified by joining a domain or a workgroup. Joining a domain involves creating an account for Oracle ZFS Storage Appliance in the given Active Directory (AD) domain. The account name can be a maximum of

15 characters, and must be unique to other names registered within the Active Directory domain. Otherwise, conflicts may occur with similarly named appliances and cause issues with functionality. After the computer account has been established, the appliance can securely query the database for information about users, groups, and shares.

Joining a workgroup implicitly leaves an Active Directory domain, and SMB clients that are stored in the Active Directory database will be unable to connect to shares.

If Active Directory is configured for the LDAP service and you are switching from domain to workgroup mode, you can optionally unconfigure or retain the LDAP configuration.

Active Directory LDAP Signing

There is no configuration option for LDAP signing, as that option is negotiated automatically when communicating with a domain controller. LDAP signing operates on communication between Oracle ZFS Storage Appliance and the domain controller, whereas SMB signing operates on communication between the SMB clients and Oracle ZFS Storage Appliance.

Related Topics

- Joining an AD Domain (BUI)
- Joining a Workgroup (BUI)
- Configuring Active Directory (CLI)
- Active Directory Join Domain
- Active Directory Windows Support

Active Directory Windows Support

Active Directory Windows Server Support

Oracle ZFS Storage Appliance software version OS8.8.x supports all Active Directory server versions.

Active Directory Windows Client Support

Starting with Windows 10, clients can enable Kerberos authentication with IP address-based Service Principal Names (SPNs) by creating a `TryIPSPN` registry entry. For more information, see Configuring Kerberos for IP Addresses. However, Oracle ZFS Storage Appliance does not support IP address-based SPNs.

Like Windows, Oracle ZFS Storage Appliance does not register IP address-based SPNs as part of the AD domain join. Microsoft notes the following potential issues with IP address-based SPNs: IP addresses are not normally used in place of hostnames because IP addresses are often temporary. Using IP addresses can lead to conflicts and authentication failures as address leases expire and renew. Therefore, registering an IP address-based SPN is a manual process and should only be used when it is not possible to switch to a DNS-based hostname.

When the `TryIPSPN` registry setting is enabled, Windows 10 clients can continue to access SMB shares exported by the appliance via `NTLMSSP`. If the domain administrator has manually added `cifs\`*`appliance-IP-address`* SPNs to the `ServicePrincipalName` attribute of the AD computer object of the appliance, Windows KDC will be able to issue CIFS service tickets for the appliance using IP address-based SPNs. However, if the client subsequently presents a CIFS service ticket that uses IP address-based SPNs to access SMB resources on the appliance, the appliance Kerberos subsystem will reject the security context due to the lack of support for IP address-based SPNs. As a result, `debug.sys` will contain the following notice-level messages:

```
krbssp: user authentication failed (GSS major error): No credentials were
supplied, or the credentials were unavailable or inaccessible
krbssp: user authentication failed (GSS minor error): No principal in keytab
('FILE:/var/krb5/krb5.keytab') matches desired name cifs/appliance-IP-address@AD-realm
```

To avoid this issue, do not publish IP-address-based SPNs to an AD server.

Related Topics

- Joining an AD Domain (BUI)
- Joining a Workgroup (BUI)
- Configuring Active Directory (CLI)
- Active Directory Join Domain
- Active Directory Domains and Workgroups

DNS Configuration

The DNS (Domain Name Service) client provides the ability to resolve IP addresses to hostnames and vice versa, and can be enabled or disabled. Optionally, secondary hostname resolution via NIS and/or LDAP, if configured and enabled, may be requested for hostnames and addresses that cannot be resolved using DNS. Hostname resolution is used throughout the appliance user interfaces, including in Logs to indicate the location from which a user performed an auditable action and in Analytics to provide statistics on a per-client basis.

To configure and manage DNS, use these procedures:

- Configuring DNS - BUI, CLI
- Using the DNS Health Tool - BUI, CLI
- Testing Hostname Resolution (CLI)
- Adding a DNS Server - BUI, CLI
- Viewing DNS Server Status - BUI, CLI

To understand DNS usage for the appliance, use these topics:

- DNS Properties and Logs
- Active Directory and DNS
- Non-DNS Resolution
- DNS-Less Operation

Configuring DNS (BUI)

DNS is usually configured during initial configuration, as described in Performing Initial Configuration (BUI) in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*. To change your DNS settings after initial configuration, use the following procedure.

1. From the **Configuration** menu, select **Services**, then **DNS**.
2. Under **General Settings**, set the following properties:
   - **DNS Domain** - Enter a domain name.
   - **DNS Search Domain(s)** - Click the add icon ⊕ to add search domain(s). To remove a domain, click the remove icon ⊖ beside it.

  - **Allow IPv4 non-DNS resolution** - Check this box to enable IPv4 non-DNS resolution. See Non-DNS Resolution.
  - **Allow IPv6 non-DNS resolution** - Check this box to enable IPv6 non-DNS resolution. See Non-DNS Resolution.

  For more information about DNS properties, see DNS Properties and Logs.

3. Click **APPLY**.

Related Topics

- Adding a DNS Server (BUI)
- DNS Health Check

Configuring DNS (CLI)

DNS is usually configured during initial configuration, as described in Performing Initial Configuration (CLI) in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*. To change your DNS settings after initial configuration, use the following procedure.

1. Go to `configuration services dns` and enter `show`.

```
hostname:> configuration services dns
hostname:configuration services dns> show
Properties:
                      <status> = online
                        domain = example.com
                       servers = IP-address
                        search =
             allow_alternate_v4 = false
             allow_alternate_v6 = false
```

2. Set the domain, servers, and search domain, and enable or disable non-DNS resolution.

   For more information, see DNS Properties and Logs and Non-DNS Resolution.

```
hostname:configuration services dns> set domain=example.com
                        domain = example.com (uncommitted)
hostname:configuration services dns> set servers=IP-address
                       servers = IP-address (uncommitted)
hostname:configuration services dns> set search=example.com
                        search = example.com (uncommitted)
hostname:configuration services dns> set allow_alternate_v4=true
            allow_alternate_v4 = true (uncommitted)
```

3. Enter `commit`.

```
hostname:configuration services dns> commit
```

Related Topics

- Adding a DNS Server (CLI)
- DNS Health Check

DNS Health Check

The DNS health tool checks the validity of your DNS configuration. If an Oracle ZFS Storage Appliance system has an issue where DNS configuration is a suspect, the DNS health tool can help identify the problem or eliminate DNS configuration as a contributor to the problem.

In addition to checking the appliance where you are running the tool, you can specify a different hostname or IP address to check.

The DNS health tool identifies questionable configuration such as the following:

- Missing or inconsistent DNS data such as missing A or PTR records or name mismatches between A and PTR records.
- Issues with looking up public Internet names.
- Issues with looking up nonexistent addresses.
- Performance problems in queries.

The DNS health tool reports the possible configuration issues that it discovers as warnings rather than as errors because that configuration might be appropriate for a specific situation. Make sure you understand the warning and the specific situation of the system that is being tested.

The following are examples of configuration issues that the DNS health tool might not identify:

- The tool cannot discover aliases, and so cannot discover issues such as the following:
  - Client references through CNAME records.
  - Names with A records that refer to the system, where those names do not appear in PTR records.
- DNS configuration that is specific to Active Directory (AD); for example, the SRV records that are used to locate domain controllers.

In addition to printing warnings, the report that is produced by the DNS health tool describes the detailed steps that were taken to discover each configuration warning. For example, if the system's configured interface addresses include an address whose PTR record refers to a different name, the report will show a description such as the following:

```
WARN: 1.2.3.4 is b.example.com, expected a.example.com
    interfaces -> 1.2.3.4
    PTR 1.2.3.4 -> b.example.com
```

You can show all lookup operations that result in a reported warning. You can show all steps performed by the DNS health tool, not just the steps that result in a reported warning.

Using the DNS Health Tool (BUI)

This procedure describes using the DNS Health Tool.

1. From the **Configuration** menu, select **Services**, then **DNS**.
2. Click the **HEALTH** tab.
3. Optional: In the **Test** section, specify the system on which to run the DNS health check.

   By default, the current system is selected. To specify a different system:

   a. Select the **Another system** radio button.

   b. In the text field, enter the hostname or IP address of the system to check.
4. Click **RUN**.
5. In the **Results** section, examine the DNS health check output.

   By default, the **WARNINGS** tab is selected, which shows only the configuration warnings.

   - To show all test steps, including steps that do not result in a warning, select the **SHOW ALL** tab.

- To show all lookup operations that result in a particular message, double-click anywhere in the message or click the information icon ⓘ to the right of the message. The additional information is shown in a separate pop-up window.

Related Topics

- Configuring DNS (BUI)
- DNS Health Check

Using the DNS Health Tool (CLI)

This procedure describes how to use the DNS Health Tool.

1. Go to `configuration services dns`.
2. Use the `health` command to run the DNS health tool on the current system or on a different system.

   The following options can be specified with the `health` command:

   `a` Show all test steps, including steps that do not result in a warning.

   `v` Show all lookup operations that result in a particular message.

   - To run the tool on this system, enter `health`.
   - To run the tool on a different system, enter one of the following commands:

     `health` *`hostname`* or `health` *`IP-address`*

Related Topics

- Configuring DNS (CLI)
- DNS Health Check

Testing Hostname Resolution (CLI)

Values returned by the CLI built-ins `nslookup` and `getent hosts` can be compared to test that hostname resolution is working:

```
hostname:> nslookup hostname
IP-address   hostname.com
hostname:> getent hosts hostname
IP-address   hostname.com
```

Adding a DNS Server (BUI)

Use the following procedure to add a DNS server.

1. From the **Configuration** menu, select **Services**, then **DNS**.
2. Click the add icon ⊕ beside **DNS Servers**.
3. In the **New DNS Server** dialog box, enter the server IP address.
4. Click **ADD**.

   A query is sent to the affected DNS servers to validate the changes. If a valid response is not received, a message appears to confirm the settings. You may confirm your changes regardless of whether the server is valid.

Related Topics

Viewing DNS Server Status (BUI)

Adding a DNS Server (CLI)

Use the following procedure to add a DNS server.

1. Go to `configuration services dns` and enter `create`.

   ```
   hostname:> configuration services dns
   hostname:configuration services dns> create
   ```

2. Enter `show`.

   ```
   hostname:configuration services server (uncommitted)> show
   Properties:
                         address = (unset)
                          status = unavailable
                             rtt = unavailable
                         err_msg =
   ```

3. Enter `set address=` and the server address.

   ```
   hostname:configuration services server (uncommitted)> set address=192.0.2.254
                        address = 192.0.2.254 (uncommitted)
   ```

4. Enter `show`.

   ```
   hostname:configuration services server (uncommitted)> show
   Properties:
                         address = 192.0.2.254
                          status = online
                             rtt = 1.812ms
                         err_msg =
   ```

5. Enter `commit`.

   A query is sent to the affected DNS servers to validate the changes. If a valid response is not received, a message appears to confirm the settings. You may confirm your changes regardless of whether the server is valid.

   ```
   hostname:configuration services server (uncommitted)> commit
   ```

Related Topics

Viewing DNS Server Status (CLI)

Viewing DNS Server Status (BUI)

Details about the DNS servers are displayed beside each entry in the BUI. A status indicator shows if the server status is online, offline, or unknown. The **RTT** column indicates the round-trip time, in milliseconds, to receive a valid response.

1. From the **Configuration** menu, select **Services**, then **DNS**.
2. Under **DNS Servers**, check the status indicator beside each server entry:

   - **Green icon** - Online
   - **Amber icon** ● - Offline
   - **Gray icon** ● - Unknown

Viewing DNS Server Status (CLI)

Select a DNS server to view its properties. The `status` property indicates if the server status is online, offline, or unknown. The `rtt` property indicates the round-trip time, in milliseconds, to receive a valid response. If the server status is offline, the `err_msg` property displays the reason, for example, `Connection timed out`.

1. Go to `configuration services dns` and enter `show`.

```
hostname:> configuration services dns
hostname:configuration services dns> show
SERVER       STATUS    ADDRESS
server-000   online    198.51.100.1
server-001   offline   198.51.100.2
```

2. Select the server for which you want to view its status.

```
hostname:configuration services dns> select server-000
```

3. Enter `show`.

```
hostname:configuration services server-000> show
Properties:
                 address = 198.51.100.1
                  status = online
                     rtt = 1.768ms
                 err_msg =
```

DNS Properties and Logs

The configurable properties for the DNS client include a base domain name and a list of servers, specified by IP address. You must supply a domain name and at least one server address; the server must be capable of returning an NS (`NameServer`) record for the domain you specify, although it need not itself be authoritative for that domain.

Table 3-5 DNS Properties

Property	Description
DNS Domain	Domain name to search first when performing partial hostname lookups.
DNS Server(s)	One or more DNS servers. IP addresses must be used.
DNS Search Domain(s)	List of up to four domains to be searched for after the Active Directory domain, the deprecated Active Directory search domain, and the specified DNS domain.
Allow IPv4 non-DNS resolution	IPv4 addresses may be resolved to hostnames, and hostnames to IPv4 addresses, using NIS and/or LDAP if configured and enabled.
Allow IPv6 non-DNS resolution	IPv4 and IPv6 addresses may be resolved to hostnames, and hostnames to IPv4 and IPv6 addresses, using NIS and/or LDAP if configured and enabled.

Changing services properties is documented in Setting Service Properties (BUI) and Setting Service Properties (CLI). The CLI property names are shorter versions of those listed above.

The DNS service events log is available in `network-dns-client:default`.

Related Topics

Active Directory and DNS

Active Directory and DNS

If you plan to use Active Directory, the servers must be able to resolve hostname and server records in the Active Directory portion of the domain namespace. For example, if your appliance resides in the domain example.com and the Active Directory portion of the namespace is `redmond.example.com`, your nameservers must be able to reach an authoritative server for `example.com`, and they must provide delegation for the domain `redmond.example.com` to one or more Active Directory servers serving that domain. These are requirements imposed by Active Directory, not Oracle ZFS Storage Appliance itself. If they are not satisfied, you will be unable to join an Active Directory domain.

> **Note:**
>
> With software version OS8.6.0 and later, if the primary DNS domain suffix does not match the DNS name of the Active Directory, the configuration results in a disjoint namespace. If you do not want a disjoint namespace, ensure that the DNS domain and the Active Directory domain are the same.

Related Topics

- DNS Configuration
- Active Directory Configuration

Non-DNS Resolution

DNS is a standard, enterprise-grade, highly scalable and reliable mechanism for mapping between hostnames and IP addresses. Use of working DNS servers is a best practice and will generally yield the best results. In some environments, there may be a subset of hosts that can be resolved only in NIS or LDAP maps. If this is the case in your environment, enable non-DNS host resolution and configure the appropriate directory service(s). If LDAP is used for host resolution, the hosts map must be located at the standard DN in your database: `ou=Hosts, (Base DN)`, and must use the standard schema. When this mode is used with NFS sharing by netgroups, it may be necessary for client systems to use the same hostname resolution mechanism configured on the appliance, or NFS sharing exceptions may not work correctly.

When non-DNS host resolution is enabled, DNS will still be used. Only if an address or hostname cannot be resolved using DNS will NIS (if enabled) and then LDAP (if enabled) be used to resolve the name or address. This can have confusing and seemingly inconsistent results. You can validate host resolution results using the `getent` CLI command described earlier.

Use of these options is strongly discouraged.

DNS-Less Operation

If the appliance is unable to access any DNS servers from its installed location in the network, you may elect to operate without DNS by supplying the server address 127.0.0.1. To operate without DNS:

- **BUI** - From the **Configuration** menu, select **Services**, then **DNS**. In the field for **DNS Server(s)**, enter **`127.0.0.1`**.

- **CLI** - Go to `configuration services dns` and enter `show`. Enter `set servers=127.0.0.1`, and then enter `commit`.

Use of this mode is strongly discouraged, because several features will not work correctly, including:

- Analytics will be unable to resolve client addresses to host names.
- The Active Directory feature will not function (you will be unable to join a domain).
- Use of SSL-protected LDAP will not work properly with certificates containing host names.
- Alert and threshold actions that involve sending e-mail can only be sent to mail servers on an attached subnet, and all addresses must be specified using the mail server's IP address.
- Some operations may take longer than normal due to hostname resolution timeouts.

These limitations may be partially mitigated by using an alternate host resolution service; see Non-DNS Resolution.

Related Topics

- Enabling a Service - BUI, CLI
- Disabling a Service - BUI, CLI

Dynamic Routing Configuration

The Routing Information Protocol (RIP) is a distance-vector dynamic routing protocol that is used by Oracle ZFS Storage Appliance to automatically configure optimal routes based on messages received from other RIP-enabled on-link hosts (typically routers). The appliance supports both RIPv1 and RIPv2 for IPv4, and RIPng for IPv6.

Routes that are configured via these protocols are marked as type `dynamic` in the routing table. RIP and RIPng listen on UDP ports 520 and 521 respectively.

Table 3-6 Dynamic Routing Log Files

Log	Description
network-routing-route:default	Logs RIP service events
network-routing-ripng:quagga	Logs RIPng service events

FTP Configuration

The FTP (File Transfer Protocol) service allows filesystem access from FTP clients. Anonymous logins are not allowed, users must authenticate with whichever name service is configured in **Services**.

FTP can be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see Kerberos Configuration. For added security when configuring FTP, you can specify the SSL/TLS versions and ciphers, as described in FTP Properties.

In a clustered environment, a share is accessible on only the controller that manages it. If the `default_root` parameter refers to a share, FTP access will be possible only from the controller that currently owns that share. If the `user_home` parameter refers to a share, automatically changing to the user's directory will be possible only from the controller that currently owns the share.

To configure FTP, use the following sections:

- Adding FTP Access to a Share (BUI)
- FTP Properties
- FTP Logs

Adding FTP Access to a Share (BUI)

Use the following procedure to add FTP access to a share.

1. From the **Configuration** menu, select **Services**.
2. Ensure that the FTP service is enabled and online. If not, enable the service.
3. Select or add a share in the **Shares** screen.
4. Click the **Protocols** tab, and check that FTP access is enabled.
5. Optional: Set the **Share mode access** to `Read only` or `Read/write`.

Related Topics

- FTP Properties
- FTP Logs

FTP Properties

The following tables describes the FTP general properties.

Table 3-7 FTP General Properties

Property	Description
Port for incoming connections	The port on which FTP listens. The default is 21.
Maximum # of connections ("0" for unlimited)	This is the maximum number of concurrent FTP connections. Set this to cover the anticipated number of concurrent users. By default this is 30, since each connection creates a system process and allowing too many (thousands) could constitute a DoS attack.
Turn on delay engine to prevent timing attacks	This inserts small delays during authentication to fool attempts at user name guessing via timing measurements. Turning this on will improve security.
Default login root	The default FTP login location that can be set so that all FTP users have a default FTP directory. • If this value is `/`, FTP users see all shares. • If this value is set to anything else (`/export` or `/export/ftp`), FTP users only see FTP shares that are under that directory. • If a valid path is provided in the User home directories field, all FTP users who have a directory under `user_home` have their default login directory set to `/export` upon login.

Table 3-7 (Cont.) FTP General Properties

Property	Description
User home directories	The location of FTP user home directories, relative to the default login root. • On login, if a user has a directory in this location, the user will be logged into that directory after successfully authenticating with the FTP service. • If the user has no home directory, the user will be logged in to the default location. Leave this property empty to disable FTP user home directories and have all users log in to the default login location.
Logging level	The verbosity of the `proftpd` log.
Permissions to mask from newly created files and directories	File permissions to remove when files are created. Group and world write are masked by default, to prevent recent uploads from being writable by everyone.

The following tables describes the FTP security properties.

Table 3-8 FTP Security Properties

Property	Description
Enable SSL/TLS	Allow SSL/TLS encrypted FTP connections. This will ensure that the FTP transaction is encrypted. The default is disabled.
SSL/TLS versions and ciphers	SSL/TLS protocol versions and ciphers for FTP connections. See Configuring SSL/TLS Versions and Ciphers.
Port for incoming SSL/TLS connections	The port that the SSL/TLS encrypted FTP service listens on. The default is 21.
Permit root login	Allow FTP logins for the root user. This is off by default, since FTP authentication is plain text which poses a security risk from network sniffing attack.
Maximum # of allowable login attempts	The number of failed login attempts before an FTP connection is disconnected, and the user must reconnect to try again. The default is 3.
Permit foreign data connection addresses	Permits foreign FTP connections to enable direct transfer of files between FTP servers. This property is off by default.

Related Topics

- Adding FTP Access to a Share (BUI)
- FTP Logs

FTP Logs

The following tables describes the FTP logs.

Table 3-9 FTP Logs

Log	Description
proftpd	Logs FTP events, including successful logins and unsuccessful login attempts.

Table 3-9 (Cont.) FTP Logs

Log	Description
proftpd_xfer	File transfer log.
proftpd_tls	Logs FTP events related to SSL/TLS encryption.

Related Topics

- Adding FTP Access to a Share (BUI)
- FTP Properties

HTTP Configuration

The HTTP service provides access to filesystems using the HTTP WebDAV (Web based Distributed Authoring and Versioning) protocol. This service allows clients to access shared filesystems through a web browser, or as a local filesystem if supported by the client software. The URL to access these HTTP shares has the following format: `http://hostname/shares/`*`mountpoint`*`/`*`share_name`*

HTTP can be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see Kerberos Configuration.

For added security when configuring HTTP, you can specify the SSL/TLS versions and ciphers, as described in HTTP Properties and Logs.

To configure HTTP, see the following sections:

- Adding HTTP Access to a Share (BUI)
- HTTP Properties and Logs
- HTTP Authentication and Access Control
- Object API Configuration

Adding HTTP Access to a Share (BUI)

Use the following procedure to add HTTP access to a share.

1. From the **Configuration** menu, select **Services**.
2. Check that the HTTP service is enabled and online. If not, enable the service.
3. Select or add a share in the **Shares** screen.
4. Click the **Protocols** tab, and check that HTTP access is enabled.
5. Optional: Set the **Share mode access** to `Read only` or `Read/write`.

   For the HTTP Object API, set the access to `Read/write`.

Related Topics

- HTTP Properties and Logs
- HTTP Authentication and Access Control
- Object API Configuration

HTTP Properties and Logs

The following tables describes the HTTP general properties.

Table 3-10 HTTP General Properties

BUI Label	CLI Property	Description
N/A	`status`	Read-only property showing the status of the HTTP service
Protocols	`protocols`	Select which access methods to support: `HTTP`, `HTTPS`, or both.
HTTP port (for incoming connections)	`listen_port`	HTTP port. The default is `80`.
HTTPS port (for incoming secure connections)	`https_port`	Secure HTTP port. The default is `443`.

The following tables describes the HTTP security properties.

Table 3-11 HTTP Security Properties

BUI Label	CLI Property	Description
SSL/TLS versions	`tls_version`	SSL/TLS protocol versions for HTTP connections. See Configuring SSL/TLS Versions and Ciphers.
List of ciphers	`ciphers`	List of SSL/TLS ciphers for HTTP connections. See Configuring SSL/TLS Versions and Ciphers.

The following tables describes the HTTP WebDAV properties.

Table 3-12 HTTP WebDAV Properties

BUI Label	CLI Property	Description
Enable WebDAV	`webdav_enabled`	When selected, enables the HTTP WebDAV feature.
Require client login	`require_login`	Clients must authenticate before share access is allowed, and files they create will have their ownership. If this property is not set, files created will be owned by the HTTP service with user `nobody`. See HTTP Authentication and Access Control.

The following tables describes the HTTP Swift object API service properties.

Table 3-13 HTTP Swift Object API Service Properties

BUI Label	CLI Property	Description
Enable Swift	`swift_enabled`	When selected, enables the HTTP Swift object API service.

Table 3-13 (Cont.) HTTP Swift Object API Service Properties

BUI Label	CLI Property	Description
Default Path	`swift_default_path`	Sets the location used when a user does not set one.

Note:

The object API service does not support changing the owner of the share. Any share owner changes will not change the account owner in the object repository and may cause subsequent authentication requests to fail.

The following tables describes the HTTP Amazon S3 object API service properties.

Table 3-14 HTTP Amazon S3 Object API Service Properties

BUI Label	CLI Property	Description
Enable S3	`s3_enabled`	When selected, enables the HTTP Amazon S3 object API service.
Default Path	`s3_default_path`	Sets the location used when a user does not set one.
Master Passphrase	`master_passphrase`	Sets the master passphrase for the Amazon S3 object API service.
Confirm Master Passphrase		Confirms the master passphrase.

Note:

The object API service does not support changing the owner of the share. Any share owner changes will not change the account owner in the object repository and may cause subsequent authentication requests to fail.

The following tables describes the HTTP Oracle Cloud Infrastructure object storage API service properties.

Table 3-15 HTTP Oracle Cloud Infrastructure Object Storage API Service Properties

BUI Label	CLI Property	Description
Enable OCI	`oci_enabled`	When selected, enables the HTTP Oracle Cloud Infrastructure Object Storage API service.
Default Path	`oci_default_path`	Sets the location used when a user does not set one.

Table 3-15 (Cont.) HTTP Oracle Cloud Infrastructure Object Storage API Service Properties

BUI Label	CLI Property	Description
Key	keys node	In the BUI, set a valid user name on the appliance, optionally enter a comment, and enter the RSA public key in PEM format for the Oracle Cloud Infrastructure Object Storage API service. In the CLI, set a valid user name on the appliance, optionally enter a comment, and use the setkey command to interactively enter the RSA public key in PEM format for the Oracle Cloud Infrastructure Object Storage API service.

> **Note:**
>
> The object API service does not support changing the owner of the share. Any share owner changes will not change the account owner in the object repository and may cause subsequent authentication requests to fail.

HTTP Logs

The HTTP service events log is available in `network-http:apache24`.

Related Topics

- Adding HTTP Access to a Share (BUI)
- HTTP Authentication and Access Control
- Object API Configuration

HTTP Authentication and Access Control

If the `Require client login` option is enabled, Oracle ZFS Storage Appliance will deny access to clients that do not supply valid authentication credentials for a local user, a NIS user, or an LDAP user. Active Directory authentication is not supported.

Only basic HTTP authentication is supported. Note that unless HTTPS is being used, this transmits the username and password unencrypted, which might not be appropriate for all environments.

Normally, authenticated users have the same permissions with HTTP that they would have with NFS or FTP. Files and directories created by an authenticated user will be owned by that user, as viewed by other protocols. Privileged users (those having a UID less than `100`) will be treated as `nobody` for the purposes of access control. Files created by privileged users will be owned by `nobody`.

If the `Require client login` option is disabled, the appliance will not try to authenticate clients (even if they do supply credentials). Newly created files are owned by `nobody`, and all users are treated as `nobody` for the purposes of access control.

Regardless of authentication, no permissions are masked from created files and directories. Created files have UNIX permissions `666` (readable and writable by everyone), and created directories have UNIX permissions `777` (readable, writable, and executable by everyone).

Related Topics

- Adding HTTP Access to a Share (BUI)
- HTTP Properties and Logs
- Object API Configuration

Object API Configuration

The object API service enables an appliance to save data as storage objects into the Oracle ZFS filesystem using the HTTP protocol and through the OpenStack Object Storage (Swift) API, the Amazon S3 (Simple Storage Service) API, or the Oracle Cloud Infrastructure Object Storage API.

After enabling and configuring the object API service, enable the feature on individual filesystems by going to the **Shares** BUI menu and selecting **Filesystems**.. Double-click on a filesystem to view its details, and then select the **Protocols** tab. In the **HTTP** section, and for the **Object store mode** option, select `Read/write` to enable the feature for the filesystem.

Enabling the Object API Service

To enable an object API service, from the **Configuration** BUI menu, select **Services**, then **HTTP**, and select **Swift**, **S3**, or **OCI**. Then, accordingly, select the check box for **Enable Swift**, **Enable S3**, or **Enable OCI**, and complete the properties using the descriptions in HTTP Properties and Logs. Finish by clicking **APPLY**.

Configuring Object API Properties

After the object API service is enabled, you can configure it with the object API properties, as shown in HTTP Properties and Logs.

Related Topics

- Adding HTTP Access to a Share (BUI)
- HTTP Properties and Logs
- HTTP Authentication and Access Control

HTTPS Configuration

The HTTPS service provides the ability to manage Oracle ZFS Storage Appliance using the HTTPS protocol. This service allows clients to manage the connection to the appliance BUI and the RESTful API service.

For added security when configuring HTTPS, you can specify the SSL/TLS versions and ciphers, and the session timeout, as described in HTTPS Properties and Logs.

HTTPS Properties and Logs

The HTTPS SSL/TLS versions and ciphers properties specify the SSL/TLS protocol versions and ciphers for HTTPS connections to Oracle ZFS Storage Appliance's BUI and the RESTful API service. See Configuring SSL/TLS Versions and Ciphers.

The HTTPS service also controls the session timeout property, which specifies the number of minutes until the browser automatically logs out of the session after user inactivity or if the user navigates away from the BUI. The default value is 15 minutes; the minimum value is 1 minute, and the maximum value is 120 minutes. This replaces the session timeout property previously located in the user preferences area of the software. Note that the CLI idle timeout property is unaffected, and it is still located under user preferences.

The HTTPS service events log is available in `appliance-kit-http:default`.

Identity Mapping Configuration

Identity mapping allows you to associate Windows and UNIX identities, thereby allowing an SMB client and an NFS client access to the same set of files. The identity mapping service manages Windows and UNIX user identities simultaneously by creating and maintaining a database of mappings between UNIX user identifiers (UIDs) and group identifiers (GIDs), and Windows security identifiers (SIDs).

To manage identity mapping, use these tasks:

- Configuring Identity Mapping - BUI, CLI
- Creating a Mapping Rule - BUI, CLI
- Viewing a Mapping (BUI)
- Flushing Mappings from the Cache - BUI, CLI

To understand identity mapping, use these topics:

- Identity Mapping Best Practices
- Identity Mapping Concepts
- Cached and Ephemeral Mappings
- Identity Mapping Case Sensitivity
- Mapping Rule Directional Symbols

Configuring Identity Mapping (BUI)

Use the following procedure to configure identity mapping.

Before You Begin

Ensure that you are joined to at least one Active Directory domain. For information about active directories, see Active Directory Configuration.

1. From the **Configuration** menu, select **Services**, then **Identity Mapping**, then **Properties**.
2. Select one of the following mapping modes.
   - **Rule-based**
   - **Directory-based** - Set all of the following attributes.

  - **AD Attribute - UNIX User Name** - Name in the Active Directory database of the equivalent UNIX user name
  - **AD Attribute - UNIX Group Name** - Name in the Active Directory database of the equivalent UNIX group name
  - **Native LDAP Attribute - Windows User Name** - Name in the LDAP database of the equivalent Windows identity
- **IDMU**

3. To save the settings, click **APPLY**. To clear the settings, click **REVERT**.

Related Topics

- For information on the different mapping modes, see Identity Mapping Concepts.
- To create an "allow" or "deny" mapping rule, see Creating a Mapping Rule (BUI).

Configuring Identity Mapping (CLI)

Use the following procedure to configure identity mapping.

Before You Begin

Ensure that you are joined to at least one Active Directory domain.

1. Go to `configuration services idmap`.
2. Enter `get` to view the identity mapping properties.

```
hostname:configuration services idmap> get

<status> = online
ad_unixuser_attr =
ad_unixgroup_attr =
nldap_winname_attr =
directory_based_mapping = none
```

```
 The three *_attr properties correspond to the three fields on C>S>Identity
Mapping>Properties.
```

3. Set `directory_based_mapping` to one of the following mapping modes.
   - To use rule-based mapping, set `directory_based_mapping` to `none`.

```
hostname:configuration services idmap> set directory_based_mapping=none
hostname:configuration services idmap>
```

   - To use directory-based mapping, set `directory_based_mapping` to `name` and assign each of the following attributes.
     - **`ad_unixuser_attr`** - Name in the Active Directory database of the equivalent UNIX user name
     - **`ad_unixgroup_attr`** - Name in the Active Directory database of the equivalent UNIX group name
     - **`nldap_winname_attr`** - Name in the LDAP database of the equivalent Windows identity

```
hostname:configuration services idmap> set directory_based_mapping=name
hostname:configuration services idmap> set ad_unixuser_attr=demo_unixuser
hostname:configuration services idmap> set ad_unixgroup_attr=demo_group
hostname:configuration services idmap> set nldap_winname_attr=demo_winuser
```

- To use Identity Management for UNIX (IDMU), set `directory_based_mapping` to `idmu`.

```
hostname:configuration services idmap> set directory_based_mapping=idmu
hostname:configuration services idmap>
```

Related Topics

- For information on the different mapping modes, see Identity Mapping Concepts.
- To create an "allow" or "deny" mapping rule, see Creating a Mapping Rule (CLI).

Creating a Mapping Rule (BUI)

Use the following procedure to grant or deny credentials for specific users through the identity mapping service. An "allow" mapping rule grants Windows identity credentials from a UNIX identity or vice versa. A "deny" mapping rule blocks a Windows identity from receiving the credentials of a UNIX identity or vice versa.

> **Note:**
>
> If you create a mapping rule that blocks a particular user, and the user's name then changes, the mapping no longer blocks that user.

Before You Begin

Configure rule-based mapping as described in Configuring Identity Mapping (BUI).

1. From the **Configuration** menu, select **Services**, then **Identity Mapping**, then **Rules**.
2. Click the add item icon ⊕ next to **Rules**.
3. In the **Add Mapping Rule** dialog box, choose either **Allow** or **Deny** for the mapping type.
4. Complete the remaining fields according to the selected mapping type.
   - **Allow mapping:**
     - **Mapping Direction** - Choose a direction.
     - **Windows Domain -** Type the Active Directory domain of the Windows identity, or select **All**.
     - **Windows Identity -** Type the name of the Windows identity.
     - **Unix Identity** - Type the name of the UNIX identity.
     - **Unix Identity Type** - Select either **User** or **Group**.
   - **Deny mapping:**
     a. For **Mapping Direction**, choose one of two options.
        - **Block Windows identity mapping** - Prevents a Windows identity from gaining the credentials of a UNIX identity.
        - **Block Windows identity mapping** - Prevents a UNIX identity from gaining the credentials of a Windows identity.
     b. Enter the Windows or UNIX identity information.
        - If you selected **Block Windows identity mapping**, type the Windows domain and identity you want to block.

  – If you selected **Block UNIX identity mapping**, type the UNIX identity and identity type you want to block.

5. Click **ADD**.

   The new mapping appears in the **Rules** list.

Related Topics

Mapping Rule Directional Symbols

Creating a Mapping Rule (CLI)

Use the following procedure to grant or deny credentials for specific users through the identity mapping service. An "allow" mapping rule grants Windows identity credentials from a UNIX identity or vice versa. A "deny" mapping rule blocks a Windows identity from receiving the credentials of a UNIX identity or vice versa.

> **Note:**
>
> If you create a mapping rule that blocks a particular user and the user's name then changes, the mapping no longer blocks that user.

Before You Begin

Configure rule-based mapping as described in Configuring Identity Mapping (CLI).

1. Go to `configuration services idmap`.
2. Enter `create`.

```
hostname:configuration services idmap>
              create

hostname:configuration services idmap (uncommitted)>
```

3. Set the properties appropriately.

   You can use the `list` command to view the available properties.

```
hostname:configuration services idmap (uncommitted)>
              list

Properties:
                    windomain = (unset)
                      winname = (unset)
                    direction = (unset)
                     unixname = (unset)
                     unixtype = (unset)
```

   a. **`windomain`** - Active Directory domain of the Windows identity.

   b. **`winname`** - Set to one of the following options.

      - To create an "allow" mapping, set `winname` to the name of the Windows identity. Enter `*` to indicate all users within the specified domain.
      - To create a "deny" mapping that blocks a UNIX identity from receiving the credentials of a Windows identity, set to the name of the Windows identity.

- To create a "deny" mapping that blocks a Windows identity from receiving the credentials of a UNIX identity, do not set `winname`.

c. **`direction`** - Set to the direction of the mapping:

- **`win2unix`** - Mapping from Windows to UNIX
- **`unix2win`** - Mapping from UNIX to Windows
- **`bi`** - Bidirectional mapping

d. **`unixname`** - Set to one of the following options:

- To create an "allow" mapping, set to the name of the UNIX identity, or enter `*` to indicate all users of the specified type.
- To create a "deny" mapping that blocks a Windows identity from receiving the credentials of a UNIX identity, set to the name of the UNIX identity.
- To create a "deny" mapping that blocks a UNIX identity from receiving the credentials of a Windows identity, do not set `unixname`.

e. **`unixtype`** - Set to either user or group for the UNIX identity type.

```
hostname:configuration services idmap (uncommitted)>
             set windomain=demo.example.com

hostname:configuration services idmap (uncommitted)>
             set winname=*

hostname:configuration services idmap (uncommitted)>
             set direction=win2unix

hostname:configuration services idmap (uncommitted)>
             set unixname=

hostname:configuration services idmap (uncommitted)>
             set unixtype=user
```

4. Enter `commit` to commit the changes, and create the mapping rule.

```
hostname:configuration services idmap (uncommitted)>
             commit

hostname:configuration services idmap>
```

You can use the `list` command to view the new rule in the Rules list.

```
hostname:configuration services idmap>
             list


MAPPING      WINDOWS ENTITY              DIRECTION        UNIX ENTITY
idmap-000    Alice@demo.example.com       (U) ==           wdp (U)
idmap-001    *@demo.example.com           (U) =>           ""  (U)
```

Example 3-1   Creating a Bi-Directional Mapping (CLI)

This example creates a bi-directional name-based mapping between a Windows user and UNIX user.

```
hostname:>
         configuration services idmap
```

```
hostname:configuration services idmap>
          create

hostname:configuration services idmap (uncommitted)>
          set
   windomain=eng.example.com

hostname:configuration services idmap (uncommitted)>
          set winname=Bill

hostname:configuration services idmap (uncommitted)>
          set direction=bi

hostname:configuration services idmap (uncommitted)>
          set unixname=wdp

hostname:configuration services idmap (uncommitted)>
          set unixtype=user

hostname:configuration services idmap (uncommitted)>
          commit

hostname:configuration services idmap>
          list

MAPPING      WINDOWS ENTITY                    DIRECTION    UNIX ENTITY
idmap-000    Bill@eng.example.com              (U) ==       wdp  (U)
```

Example 3-2   Creating a Deny Mapping (CLI)

This example creates a deny mapping to prevent all Windows users in a domain from obtaining credentials.

```
hostname:configuration services idmap>
          create

hostname:configuration services idmap (uncommitted)>
          list

Properties:
                     windomain = (unset)
                       winname = (unset)
                     direction = (unset)
                      unixname = (unset)
                      unixtype = (unset)

hostname:configuration services idmap (uncommitted)>
          set

   windomain=guest.example.com
hostname:configuration services idmap (uncommitted)>
          set winname=*

hostname:configuration services idmap (uncommitted)>
          set direction=win2unix

hostname:configuration services idmap (uncommitted)>
          set unixname=

hostname:configuration services idmap (uncommitted)>
```

```
          set unixtype=user

hostname:configuration services idmap (uncommitted)>
          commit

hostname:configuration services idmap>
          list

MAPPING      WINDOWS ENTITY                  DIRECTION   UNIX ENTITY
idmap-000    Bill@eng.example.com            (U) ==      wdp  (U)
idmap-001    *@guest.example.com             (U) =>      ""   (U)
```

Viewing a Mapping (BUI)

Use the following procedure to view an existing mapping.

1. From the **Configuration** menu, select **Services**, then **Identity Mapping**, then **Show Mappings**.
2. Choose either **Windows** or **UNIX** for the platform from which the identity is mapped.
3. Enter the Windows or UNIX identity information.
   - If you selected **Windows**, type the Windows domain and name of the user.
   - If you selected **UNIX**, choose either **User** or **Group** for the type, and type the entity name.
4. Click **SHOW MAPPING**.

   The identity user or group properties are displayed. The mapping source and backend origin are also displayed:

   Source:

   - **New mapping** - The mapping was newly created and was neither retrieved from the cache nor predefined.
   - **Cached mapping** - The mapping was retrieved from the cache, where mappings are stored for 10 minutes after they are requested.
   - **Hard coded mapping** - The mapping is predefined and fixed on the appliance. These mappings were created for default UNIX and Windows identities.
   - **Algorithmic mapping** - A non-ephemeral UNIX UID or GID could not be mapped by name, so it was mapped to an algorithmically generated SID.

   Backend:

   - **AD Directory** - This is a directory-based mapping that was created using annotations in the Active Directory.
   - **Native LDAP Directory** - This is a directory-based mapping that was created using annotations in the LDAP directory.
   - **IDMU** - The mapping was created using the Windows feature Identity Management for UNIX.
   - **Name rule** - The mapping was created using a name rule.
   - **Ephemeral** - Since there was no equivalent identity at the time the mapping was created, the system created a temporary one using an ephemeral UID or GID.
   - **Local SID** - A non-ephemeral UNIX UID or GID could not be mapped by name, so it was mapped to an algorithmically generated local SID.

- **Well-known mapping** - The mapping uses a "well-known SID." These Windows SIDs identify generic users or generic groups. Their values remain constant across all operating systems.

Flushing Mappings from the Cache (BUI)

Use the following procedure to flush, or expire, all mappings from the cache.

After a requested mapping has been provided, it is stored in the cache for 10 minutes and then expires. You can immediately expire a mapping by using the flush function, which expires all cached mappings.

1. From the **Configuration** menu, select **Services** , then **Identity Mapping**, then **Show Mappings**.
2. Click **FLUSH MAP CACHE**.

   All cached mappings are expired.

Flushing Mappings from the Cache (CLI)

Use the following procedure to flush, or expire, all mappings from the cache.

After a requested mapping has been provided, it is stored in the cache for 10 minutes and then expires. You can immediately expire a mapping by using the flush function, which expires all cached mappings.

1. Go to `configuration services idmap`.
2. Enter `flush`.

   ```
   hostname:configuration services idmap> flush
   hostname:configuration services idmap>
   ```

   All cached mappings are expired.

Identity Mapping Best Practices

- Configure user-specific identity mapping rules when you want a user to have access to a common set of files through both NFS and SMB clients. If NFS and SMB clients are accessing disjointed filesystems, there is no need to configure any identity mapping rules.
- Reconfiguring the identity mapping service does not affect active SMB sessions. Connected users remain connected, and their previous name mapping is available for authorizing access to additional shares for up to 10 minutes. To prevent unauthorized access, configure the mappings before exporting shares.
- The security that your identity mappings provide is only as good as their synchronization with your directory services. For example, if you create a name-based mapping that denies access to a particular user, and the user's name changes, the mapping no longer denies access to that user.
- You can only have one bidirectional mapping for each Windows domain that maps all users in the Windows domain to all UNIX identities. If you want to create multiple domain-wide rules, be sure to specify that those rules map *only* from Windows to UNIX.
- Use the IDMU mapping mode instead of directory-based mapping whenever possible.

Identity Mapping Concepts

The SMB service uses the identity mapping service to associate Windows and UNIX identities. When the SMB service authenticates a user, it uses the identity mapping service to map the user's Windows identity to the appropriate UNIX identity. If no UNIX identity exists for a Windows user, the service generates a temporary identity using an ephemeral UID and GID. These mappings allow a share to be exported and accessed concurrently by SMB and NFS clients. By associating Windows and UNIX identities, NFS and SMB clients can share the same identity, thereby allowing access to the same set of files.

In the Windows operating system, an access token contains the security information for a login session and identifies the user, the user's groups, and the user's privileges. Administrators define Windows users and groups in a Workgroup, or in a SAM database, which is managed on an Active Directory domain controller. Each user and group has a SID, which uniquely identifies the user or group, both within a host and a local domain, and across all possible Windows domains.

UNIX creates user credentials based on user authentication and file permissions. Administrators define UNIX users and groups in local password and group files or in a name or directory service, such as NIS or LDAP. Each UNIX user and group has a UID and GID. Typically, the UID or GID uniquely identifies a user or group within a single UNIX domain. However, these values are not unique across domains.

The following options are available when selecting a mapping mode:

- **Rule-based Mapping** - Use for creating various rules that map identities by name, thus establishing equivalences between Windows and UNIX identities. Mapping rules are useful when you want a user to access the same set of files through both SMB and NFS clients.
- **Directory-based Mapping** - Use for annotating an LDAP or Active Directory object with information about how the identity maps to an equivalent identity on the opposite platform.
- **IDMU-based Mapping** - Identity Management for UNIX (IDMU) is a feature that Microsoft offers for Windows Server 2003, and is bundled with Windows Server 2003 R2 and later. IDMU supports Windows as a NIS/NFS server by adding a "UNIX Attributes" panel to the Active Directory Users and Computers user interface. This allows administrators to specify a number of UNIX-related parameters, including UID, GID, login shell, and home directory. These parameters are made available through Active Directory using a schema similar to, but not the same as, RFC 2307, and through the NIS service. When the IDMU mapping mode is selected, the identity mapping service consumes these UNIX attributes to establish mappings between Windows and UNIX identities. This approach is very similar to directory-based mapping, except that the identity mapping service queries the property schema established by the IDMU software instead of allowing a custom schema. When this approach is used, no other directory-based mapping may occur.

Cached and Ephemeral Mappings

When the identity mapping service provides a name mapping, it stores the mapping in the cache for 10 minutes, at which point the mapping expires. Within its 10-minute life, a mapping is persistent across restarts of the identity mapping service. Changes to the mappings or to the name service directories do not affect existing connections within the 10-minute life of a mapping. The service evaluates mappings only when the client tries to connect to a share and there is no unexpired mapping. For example, if the SMB server requests a mapping for the user after the mapping has expired, the service re-evaluates the mapping.

If no name-based mapping rule applies for a particular user, that user will be given temporary credentials through an ephemeral mapping unless the user is blocked by another mapping.

When a Windows user with an ephemeral UNIX name creates a file on the system, Windows clients accessing the file using SMB see that the file is owned by that Windows identity. However, NFS clients see that the file is owned by "nobody".

Identity Mapping Case Sensitivity

Windows names are not case sensitive, but UNIX names are case sensitive. The user names `JSMITH`, `JSmith`, and `jsmith` are equivalent names in Windows, but they are three distinct names in UNIX. Case sensitivity affects name mappings differently depending on the direction of the mapping.

- For a Windows-to-UNIX mapping to produce a match, the case of the Windows user name must match the case of the UNIX user name. For example, only Windows user name `jsmith` matches UNIX user name `jsmith`. Windows user name `Jsmith` does not match.
- An exception to the case matching requirement for Windows-to-UNIX mappings occurs when the mapping uses the wildcard character "*" to map multiple user names.

  If the identity mapping service encounters a mapping that maps Windows user `*@some.domain to UNIX user "*"`, it first searches for a UNIX name that matches the Windows name exactly. If it does not find a match, the service converts the entire Windows name to lower case and searches again for a matching UNIX name. For example, the Windows user name `JSmith@some.domain` maps to UNIX user name `jsmith`. If the service does not find a match after using lowercase for the Windows user name, the user does not obtain a mapping.

  You can create a rule to match strings that differ only in case. For example, you can create a user-specific mapping to map the Windows user `JSmith@some.domain` to UNIX user `jSmith`. Otherwise, the service assigns an ephemeral ID to the Windows user.
- For a UNIX-to-Windows mapping to produce a match, the case does not have to match. For example, UNIX user name `jsmith` matches any Windows user name with the letters `JSMITH` regardless of case.

Mapping Rule Directional Symbols

After creating a name-based mapping, the following symbols indicate the semantics of each rule.

- ⇔ - Maps Windows identity to UNIX identity and UNIX identity to Windows identity
- o→ - Maps Windows identity to UNIX identity
- ←o - Maps UNIX identity to Windows identity
- o⊣ - Prevents Windows identity from obtaining credentials
- ⊢o - Prevents UNIX identity from obtaining credentials

If an icon is gray instead of black, the rule matches a UNIX identity that cannot be resolved.

IPMP Configuration

Internet Protocol Network Multipathing (IPMP) allows multiple network interfaces to be grouped as one, both for improved network bandwidth and for reliability (interface redundancy). Some properties can be configured in this section. For the configuration of network interfaces in IPMP groups, see Network Configuration.

Table 3-16 IPMP Properties

Property	Description
Failure detection latency	Time for IPMP to declare a network interface has failed, and to fail over its IP addresses.
Enable fail-back	Allow the service to resume connections to a repaired interface.

Changing services properties is documented in Setting Service Properties (BUI) and Setting Service Properties (CLI). The CLI property names are shorter versions of those listed earlier.

The IPMP service events log is available in `network-initial:default`.

iSCSI Configuration

When you configure a LUN on the appliance you can export that volume over an Internet Small Computer System Interface (iSCSI) target. The iSCSI service allows iSCSI initiators to access targets using the iSCSI protocol.

The service supports discovery, management, and configuration using the iSNS protocol. The iSCSI service supports both unidirectional (target authenticates initiator) and bidirectional (target and initiator authenticate each other) authentication using CHAP. Additionally, the service supports CHAP authentication data management in a RADIUS database.

The system performs authentication first, and authorization second, in two independent steps.

> **Note:**
>
> For examples of configuring iSCSI initiators and targets, see Configuring Storage Area Network (SAN).

Table 3-17 iSCSI Service Properties

Property	Description
Use iSNS	Whether iSNS discovery is enabled
iSNS Server	An iSNS server
Use RADIUS	Whether RADIUS is enabled
RADIUS Server	A RADIUS server
RADIUS Server Secret	The RADIUS server's secret

If the local initiator has a CHAP name and a CHAP secret, the system performs authentication. If the local initiator does not have the CHAP properties, the system does not perform any authentication and therefore all initiators are eligible for authorization.

The iSCSI service allows you to specify a global list of initiators that you can use within initiator groups.

If your initiator cannot connect to your target:

- Make sure the IQN of the initiator matches the IQN identified in the initiators list.
- Check that IP address of iSNS server is correct and that the iSNS server is configured.

- Check that the IP address of the target is correct on the initiator side.
- Check that initiator CHAP names and secrets match on both sides.
- Make sure that the target CHAP name and secret do not match those of any of the initiators.
- Check that the IP address and secret of the RADIUS server are correct, and that the RADIUS server is configured.
- Check that the initiator accessing the LUN is a member of that LUN's initiator group.
- Check that the targets exporting that LUN are online.
- Check that the LUN's operational status is online.
- Check the logical unit number for each LUN.

If, during the failover / failbacks, the iSER `Reduced Copy I/Os` from the Red Hat client are not surviving, modify the `node.session.timeo.replacement_timeout` parameter in the `/etc/iscsi/iscsid.conf` file to 300sec.

Related Topics

Setting Service Properties - BUI, CLI

Kerberos Configuration

Kerberos is a network protocol that uses secret-key cryptography to authenticate communication between a client and a host machine or service. It uses a Key Distribution Center (KDC) server to issue time-stamped tickets. You can use the appliance to import Kerberos principals and keys created on the KDC, or you can configure principals for the KDC using the appliance, and their keys are automatically created. Although you can use both methods, importing is the best practice and most commonly used. All keys are encrypted using the Kerberos password and stored within the appliance keytab file.

Both Kerberos and Active Directory can be enabled at the same time because they have distinct realms and keys. When both are active, the Kerberos realm is the default.

The appliance can use Kerberos to authenticate users for administrative login and for access to services, including NFS, HTTP, FTP, SFTP, and SSH. An appliance user must have a Kerberos principal by the same name to use Kerberos authentication for these services. Kerberos can also be used to set security for individual shares that use the NFS protocol, as described in Configuring Kerberos Realms for NFS. Since the Kerberos service uses time stamps, configure the appliance NTP service first.

To configure Kerberos, see the following sections:

- Creating a Kerberos Realm - BUI, CLI
- Importing Kerberos Keys - BUI, CLI
- Creating Kerberos Principals and Keys - BUI, CLI
- Deleting Kerberos Principals and Keys - BUI, CLI
- Destroying a Kerberos Realm - BUI, CLI
- Kerberos Service Properties
- Kerberos Properties and Logs
- Configuring Kerberos Realms for NFS

Creating a Kerberos Realm (BUI)

Use the following procedure to create a Kerberos realm, set the KDC(s), and select strong or weak encryption types. Descriptions of each property are located in Kerberos Service Properties.

Before You Begin

Ensure that you have configured the NTP service.

1. From the **Configuration** menu, select **Services**.
2. To enable the Kerberos service, click the enable icon ⏻ for **Kerberos**.
3. Click **Kerberos**.
4. In the **Realm** field, type the Kerberos realm.

   For familiarity, the realm name can be the same as your DNS domain name, except that the realm name is in uppercase.

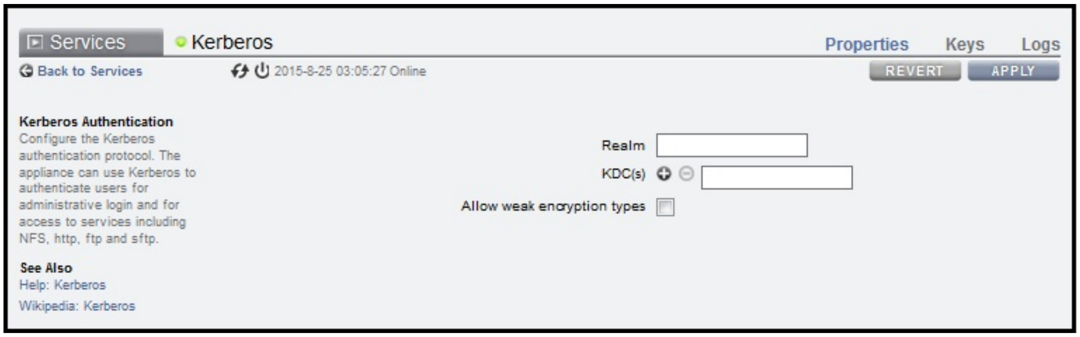

Kerberos Authentication
Configure the Kerberos authentication protocol. The appliance can use Kerberos to authenticate users for administrative login and for access to services including NFS, http, ftp and sftp.

See Also
Help: Kerberos
Wikipedia: Kerberos

Realm
KDC(s) ⊕ ⊖
Allow weak encryption types □

5. In the KDC(s) field, type the host name of the KDC administrative server.

   If your Kerberos configuration includes DNS support for KDC lookup, leave this field blank.
6. If you have another KDC, click the add icon ⊕ next to KDC(s) and type its host name. Repeat for each additional KDC.

   If your configuration includes DNS support, do not complete this step.
7. To allow support for weak encryption types, such as DES and Exportable ArcFour with HMAC/md5, select **Allow weak encryption types**.

   The default does not support weak encryption types.
8. Click **APPLY**.

   To reset the properties to their original values, click **REVERT** instead.

Next Steps

Choose one of the following options:

- Importing Kerberos Keys (BUI)
- Creating Kerberos Principals and Keys (BUI)

Creating a Kerberos Realm (CLI)

Use the following procedure to create a Kerberos realm, set the KDC(s), and select strong or weak encryption types. Descriptions of each property are located in Kerberos Service Properties and Kerberos Properties and Logs.

Before You Begin

Ensure that you have configured the NTP service.

1. Go to `configuration services kerberos` and enter `show`.

```
hostname:configuration services kerberos>
              show

Properties:
                     <status> = disabled
            allow_weak_crypto = false
```

2. To enable the Kerberos service, enter `enable` and then enter `commit`.
3. To allow support for weak encryption types, such as DES and Exportable ArcFour with HMAC/md5, enter `set allow_weak_crypto=true` and then enter `commit`.

   The default does not support weak encryption types.
4. To create a realm, enter `create` and the realm name, and then enter `commit`.

   For familiarity, the realm name can be the same as your DNS domain name, except that the realm name is in uppercase.

```
hostname:configuration services kerberos>
              create TEST.NET

hostname:configuration services kerberos TEST.NET (uncommitted)>
              commit
```

5. Enter `done`.
6. To view all realms, enter `list`.

```
hostname:configuration services kerberos>
              list

REALM               KDC
TEST.NET
```

7. Select the realm.

```
hostname:configuration services kerberos>
              select TEST.NET

hostname:configuration services kerberos TEST.NET>
```

8. To configure the KDC server(s), enter `set kdcs=` and the KDC administrative server host name. If you have additional KDCs, add them to the same line and separate them by commas. Then enter `commit`.

   If your Kerberos configuration includes DNS support for KDC lookup, do not perform this step.

```
hostname:configuration services kerberos TEST.NET>
                    set kdcs=kdc1.example.com,kdc2.example.com

                     kdcs = kdc1.example.com,kdc2.example.com (uncommitted)
hostname:configuration services kerberos TEST.NET>
                    commit
```

Next Steps

Choose one of the following options:

- Importing Kerberos Keys (CLI)
- Creating Kerberos Principals and Keys (CLI)

Importing Kerberos Keys (BUI)

Use the following procedure to import Kerberos keys that were created on the KDC. The keys are then stored in the appliance keytab. This task does not require login credentials on the KDC. Descriptions of each property are located in Kerberos Service Properties.

Before You Begin

Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in Creating a Kerberos Realm (BUI).

1. From the **Configuration** menu, select **Services**.
2. Click **Kerberos**.
3. Click **Keys** and click **IMPORT KEYS**.

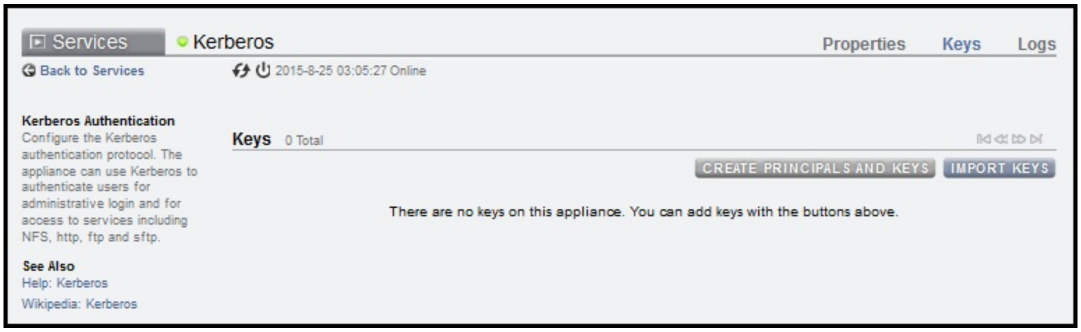

Configure the Kerberos authentication protocol. The appliance can use Kerberos to authenticate users for administrative login and for access to services including NFS, http, ftp and sftp.

There are no keys on this appliance. You can add keys with the buttons above.

4. In the **Import Keys** dialog box, click **Browse** and select the Kerberos keytab file.
5. Click **UPLOAD**.

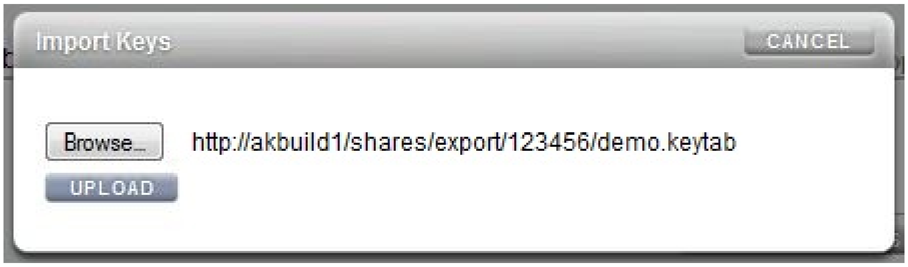

http://akbuild1/shares/export/123456/demo.keytab

The list of keys is displayed.

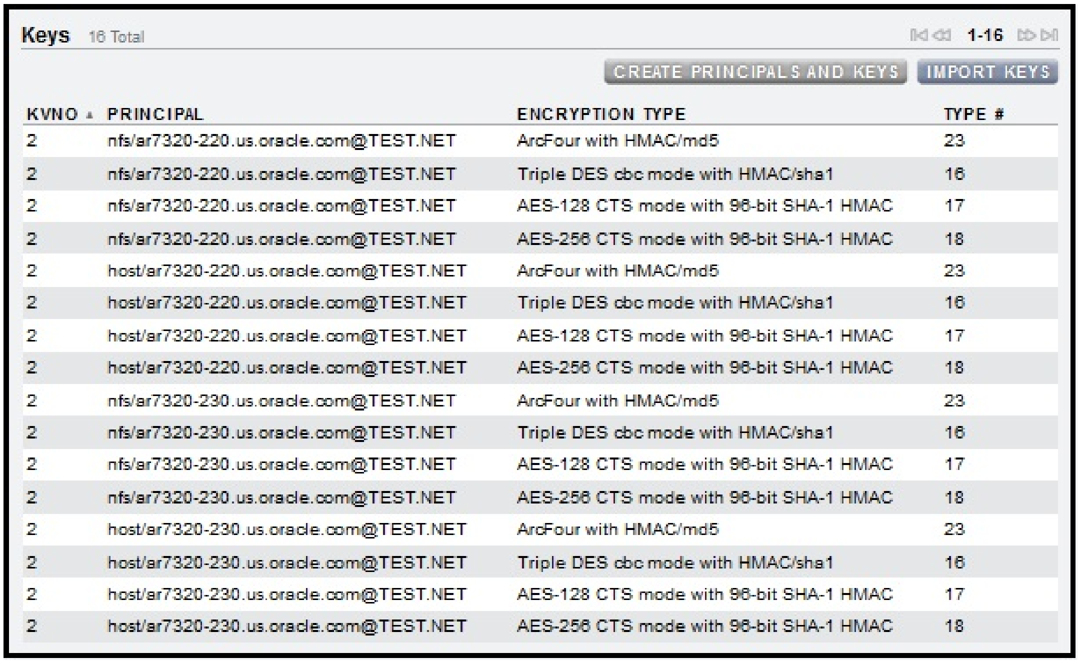

KVNO ▴	PRINCIPAL	ENCRYPTION TYPE	TYPE #
2	nfs/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	nfs/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18

Importing Kerberos Keys (CLI)

Use the following procedure to import Kerberos keys that were created on the KDC. The keys are then stored in the appliance keytab. This task does not require login credentials on the KDC. Descriptions of each property are located in Kerberos Service Properties and Kerberos Properties and Logs.

Before You Begin

Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in Creating a Kerberos Realm (CLI).

1. Go to `configuration services kerberos importkeytab` and enter `show` to view the properties.

   ```
   hostname:configuration services kerberos importkeytab (uncommitted)>
                show

   Properties:
                            url = (unset)
                           user = (unset)
                       password = (unset)
   ```

2. Enter `set url=` and the URL of the Kerberos keytab file.

   ```
   hostname:configuration services kerberos importkeytab (uncommitted)>
                set url=http://akbuild1/shares/export/123456/demo.keytab

                            url = http://akbuild1/shares/export/123456/demo.keytab
   ```

3. Enter `set user=` and the user name for URL access.

   ```
   hostname:configuration services kerberos importkeytab (uncommitted)>
                set user=myusername
   ```

```
                         user = myusername
```

4. Enter set `password=` and the password for URL access, and then enter `commit`.

   ```
   hostname:configuration services kerberos importkeytab (uncommitted)>
                 set password=letmein

                        password = (set)
   hostname:configuration services kerberos importkeytab (uncommitted)>
                 commit

   Transferred 718 of 718 (100%) . . . done
   Imported 8 keys.
   ```

5. Enter `show` to view the realms and KDCs.

   ```
   hostname:configuration services kerberos>
                 show

   Properties:
                        <status> = online
               allow_weak_crypto = true
   Realms:
   REALM          KDC
   TEST.NET       kdc1.example.com
   ```

6. To view the principals for a realm, select a realm and enter `show`.

   ```
   hostname:configuration services kerberos>
                 select TEST.NET

   hostname:configuration services kerberos TEST.NET>
                 show

   Properties:
                   kdcs = kdc1.example.com
   Keytab entries:
   NAME             KEYS  PRINCIPAL
   principal-000    4     host/hostname.example.com@TEST.NET
   principal-001    4     nfs/hostname.example.com@TEST.NET
   ```

7. To view the keys for a principal, select a principal and enter `show`.

   ```
   hostname:configuration services kerberos TEST.NET>
                 select principal-001

   hostname:configuration services kerberos principal-001>
                 show

   Properties:
                    name = nfs/hostname.example.com@TEST.NET
   Keys:
   KEY        KVNO   ENCTYPENO   ENCTYPE
   key-000    28     18          AES-256 CTS mode with 96-bit SHA-1 HMAC
   key-001    28     17          AES-128 CTS mode with 96-bit SHA-1 HMAC
   key-002    28     16          Triple DES cbc mode with HMAC/sha1
   key-003    28     23          ArcFour with HMAC/md5
   key-004    28     24          Exportable ArcFour with HMAC/md5
   key-005    28     3           DES cbc mode with RSA-MD5
   key-006    28     1           DES cbc mode with CRC-32
   ```

   Legend for column headings:

- KEY = Key name
- KVNO = Key version number
- ENCTYPENO = Encryption type number
- ENCTYPE = Encryption type

8. To view the properties of a key, select a key and enter show.

```
hostname:configuration services kerberos principal-001>
              select key-003

hostname:configuration services kerberos principal-001 key-003>
              show

Properties:
               principal = nfs/hostname.example.com@TEST.NET
                    kvno = 28
                 enctype = ArcFour with HMAC/md5
               enctypeno = 23
```

Creating Kerberos Principals and Keys (BUI)

Use the following procedure to create Kerberos principals on the KDC administrative server using the appliance. Keys are generated for each principal and stored in the appliance keytab. Descriptions of each property are located in Kerberos Service Properties.

Before You Begin

- Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in Creating a Kerberos Realm (BUI).
- Ensure that you have login credentials on the KDC.

1. From the **Configuration** menu, select **Services**.
2. Click **Kerberos**.
3. Click **Keys** and click **CREATE PRINCIPALS AND KEYS**.

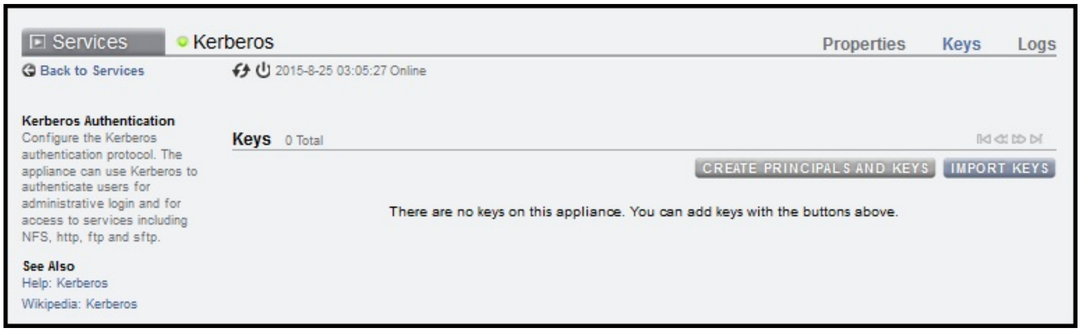

Kerberos Authentication

Configure the Kerberos authentication protocol. The appliance can use Kerberos to authenticate users for administrative login and for access to services including NFS, http, ftp and sftp.

There are no keys on this appliance. You can add keys with the buttons above.

4. In the **KDC Admin Login** dialog box, complete the following fields:
   - **Realm -** This field is auto-populated and cannot be modified.
   - **Admin server -** KDC administrative server host name. This field is auto-populated, but can be modified.
   - **Admin principal -** KDC administrator name for the realm.
   - **Password -** Password for the KDC administrator.

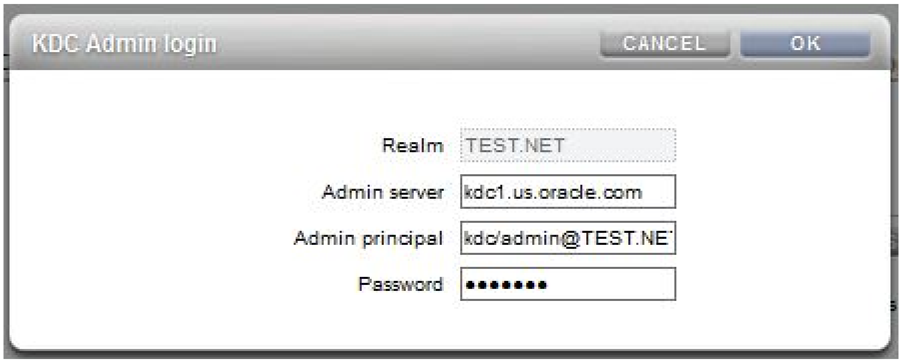

Realm: TEST.NET
Admin server: kdc1.us.oracle.com
Admin principal: kdc/admin@TEST.NE
Password: •••••••

5. Click **OK**.
6. In the confirmation box, click **OK**.

   The list of principals and keys is displayed.

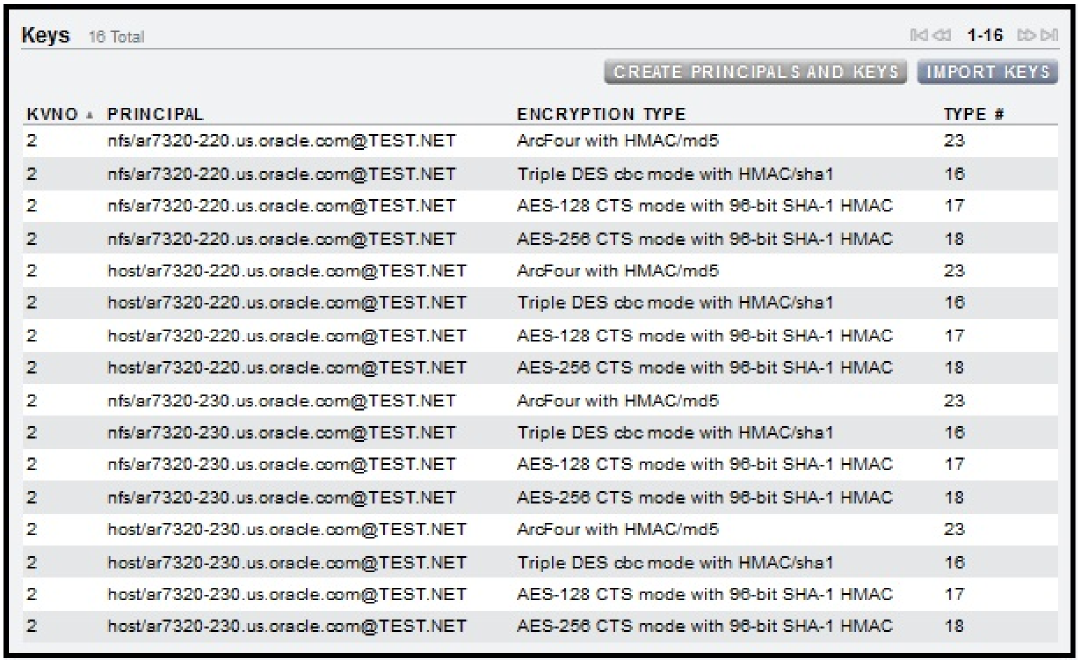

KVNO ▲	PRINCIPAL	ENCRYPTION TYPE	TYPE #
2	nfs/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	nfs/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18

Creating Kerberos Principals and Keys (CLI)

Use the following procedure to create Kerberos principals on the KDC administrative server using the appliance. Keys are generated for each principal and stored in the appliance keytab. Descriptions of each property are located in Kerberos Service Properties and Kerberos Properties and Logs.

Before You Begin

- Ensure that you have enabled the Kerberos service, set the realm, and identified the KDC(s) as described in Creating a Kerberos Realm (CLI).
- Ensure that you have login credentials on the KDC.

1. Go to `configuration services kerberos` and enter `list`.

```
hostname:configuration services kerberos>
              list

REALM               KDC
TEST.NET
```

2. Select the realm.

```
hostname:configuration services kerberos>
              select TEST.NET

hostname:configuration services kerberos TEST.NET>
```

3. To create the principals, enter principals and then enter `show` to view the properties.

```
hostname:configuration services kerberos TEST.NET>
              principals

hostname:configuration services kerberos TEST.NET principals (uncommitted)>
              show

Properties:
               realm = TEST.NET
              server = kdc1.example.com
               admin = (unset)
            password = (unset)
```

4. Optional: To change the KDC server, enter `set kdcs=` and the KDC server host name. Then enter commit.

```
hostname:configuration services kerberos TEST.NET>
              set kdcs=kdc2.example.com

               kdcs = kdc2.example.com (uncommitted)
hostname:configuration services kerberos TEST.NET>
              commit
```

5. Enter `set admin=` and the KDC administrator name for the realm.

```
hostname:configuration services kerberos TEST.NET principals (uncommitted)>
              set admin=kdc/admin
```

6. Enter `set password=` and the KDC administrator password, and then enter `commit`.

```
hostname:configuration services kerberos TEST.NET principals (uncommitted)>
              set password=test123

             password = (set)
hostname:configuration services kerberos TEST.NET principals (uncommitted)>
              commit
```

7. Enter `show` to view the principals for the KDC.

```
hostname:configuration services kerberos TEST.NET>
              show

Properties:
                kdcs = kdc1.example.com
Keytab entries:
```

```
NAME            KEYS  PRINCIPAL
principal-000   4     host/hostname.example.com@TEST.NET
principal-001   4     nfs/hostname.example.com@TEST.NET
```

8. To view the keys for a principal, select a principal and enter `show`.

```
hostname:configuration services kerberos TEST.NET>
              select principal-001

hostname:configuration services kerberos principal-001>
              show

Properties:
                 name = nfs/hostname.example.com@TEST.NET
Keys:
KEY       KVNO   ENCTYPENO   ENCTYPE
key-000   28     18          AES-256 CTS mode with 96-bit SHA-1 HMAC
key-001   28     17          AES-128 CTS mode with 96-bit SHA-1 HMAC
key-002   28     16          Triple DES cbc mode with HMAC/sha1
key-003   28     23          ArcFour with HMAC/md5
key-004   28     24          Exportable ArcFour with HMAC/md5
key-005   28     3           DES cbc mode with RSA-MD5
key-006   28     1           DES cbc mode with CRC-32
```

Legend for column headings:

- `KEY` = Key name
- `KVNO` = Key version number
- `ENCTYPENO` = Encryption type number
- `ENCTYPE` = Encryption type

9. To view the properties of a key, select a key and enter `show`.

```
hostname:configuration services kerberos principal-001>
              select key-003

hostname:configuration services kerberos principal-001 key-003>
              show

Properties:
               principal = nfs/hostname.example.com@TEST.NET
                    kvno = 28
                 enctype = ArcFour with HMAC/md5
               enctypeno = 23
```

Deleting Kerberos Principals and Keys (BUI)

Use the following procedure to delete individual keys.

1. From the **Configuration** menu, select **Services**.
2. Click **Kerberos**.
3. Click **Keys**.

   The list of principals and keys is displayed.

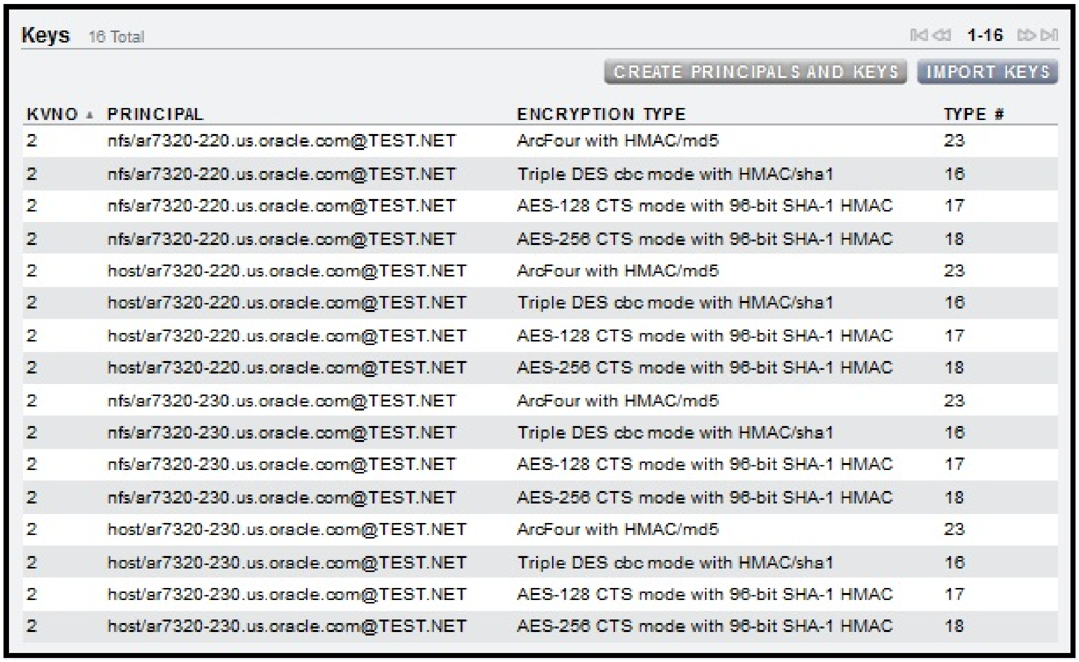

KVNO ▲	PRINCIPAL	ENCRYPTION TYPE	TYPE #
2	nfs/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-220.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-220.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-220.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	nfs/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	nfs/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	nfs/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18
2	host/ar7320-230.us.oracle.com@TEST.NET	ArcFour with HMAC/md5	23
2	host/ar7320-230.us.oracle.com@TEST.NET	Triple DES cbc mode with HMAC/sha1	16
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-128 CTS mode with 96-bit SHA-1 HMAC	17
2	host/ar7320-230.us.oracle.com@TEST.NET	AES-256 CTS mode with 96-bit SHA-1 HMAC	18

4. Optional: To sort by a column, such as `PRINCIPAL`, click the column heading.
5. To delete an individual key, hover over the appropriate row, click its trash icon 🗑 and confirm your action.

   If you delete all keys for a principal, you effectively delete the principal from the appliance.

Deleting Kerberos Principals and Keys (CLI)

Use the following procedure to delete individual keys, or to delete all keys for a principal.

1. Go to `configuration services kerberos` and enter `list`.

```
hostname:configuration services kerberos>
              list

REALM               KDC
TEST.NET
```

2. Select the realm.

```
hostname:configuration services kerberos>
              select TEST.NET

hostname:configuration services kerberos TEST.NET>
```

3. Enter `show` to view the principals for the KDC.

```
hostname:configuration services kerberos TEST.NET>
              show

Properties:
                kdcs = kdc1.example.com
Keytab entries:
NAME            KEYS   PRINCIPAL
principal-000   4      host/hostname.example.com@TEST.NET
principal-001   4      nfs/hostname.example.com@TEST.NET
```

4. To delete all of the keys for a principal, enter `destroy` and the principal name, and confirm your action.

   To delete an individual key, see the next step.

```
hostname:configuration services kerberos TEST.NET>
              destroy principal-000

This will delete all keys for "principal-000". Are you sure? (Y/N)
              Y
```

5. To delete an individual key for a principal, first select a principal and enter `show` to view the list of keys.

```
hostname:configuration services kerberos TEST.NET>
              select principal-001

hostname:configuration services kerberos principal-001>
              show

Properties:
                 name = nfs/hostname.example.com@TEST.NET
Keys:
KEY       KVNO   ENCTYPENO   ENCTYPE
key-000   28     18          AES-256 CTS mode with 96-bit SHA-1 HMAC
key-001   28     17          AES-128 CTS mode with 96-bit SHA-1 HMAC
key-002   28     16          Triple DES cbc mode with HMAC/sha1
key-003   28     23          ArcFour with HMAC/md5
key-004   28     24          Exportable ArcFour with HMAC/md5
key-005   28     3           DES cbc mode with RSA-MD5
key-006   28     1           DES cbc mode with CRC-32
```

   Legend for column headings:

   - `KEY` = Key name
   - `KVNO` = Key version number
   - `ENCTYPENO` = Encryption type number
   - `ENCTYPE` = Encryption type

6. To view the properties of a key, select a key and enter `show`.

```
hostname:configuration services kerberos principal-001>
              select key-003

hostname:configuration services kerberos principal-001 key-003>
              show

Properties:
               principal = nfs/hostname.example.com@TEST.NET
                    kvno = 28
                 enctype = ArcFour with HMAC/md5
               enctypeno = 23
```

7. To delete a key or view a different key, enter `done` to return to the `principal` context.

```
hostname:configuration services kerberos principal-001 key-003>
              done

hostname:configuration services kerberos principal-001>
```

8. To delete the key, enter `destroy` and the key name, and confirm your action.

```
hostname:configuration services kerberos principal-001>
              destroy key-003

This will delete key "key-003". Are you sure? (Y/N)
              Y
```

Destroying a Kerberos Realm (BUI)

Destroying a realm also destroys its corresponding keys.

1. From the **Configuration** menu, select **Services**.
2. Click **Kerberos**.
3. Clear the **Realm** field, click **APPLY**, and confirm your action.

Destroying a Kerberos Realm (CLI)

Destroying a realm also destroys its corresponding keys.

1. Go to `configuration services kerberos`.
2. Enter `destroy` and the realm name, and then confirm your action.

```
hostname:configuration services kerberos> destroy TEST.NET
This will destroy "TEST.NET". Are you sure? (Y/N) Y
```

Kerberos Service Properties

The following properties are available for the Kerberos service:

- **Realm -** A string, the name of the realm.
- **KDC(s) -** A list of zero or more host names, the Key Distribution Center(s) for the realm. The first Key Distribution Center (KDC) listed is assumed to be the Admin Server, which is relevant if creating principals on the KDC from the appliance, but not when importing keys. The list may be empty if at least one KDC is published for the realm in DNS.
- **Allow weak encryption types -** A Boolean value. This enables/disables support for deprecated weak encryption types (des-cbc-crc, des-cbc-md5, and arcfour-hmac-exp). This property is disabled by default.
- **Admin -** A string, the name of the Kerberos admin principal (administrator). By convention, a principal name is divided into three components: the primary, the instance, and the realm. You can specify a principal as `joe`, `joe/admin`, or `joe/admin@ENG.EXAMPLE.COM`. This property is used only if creating service principals, and is not retained.
- **Password -** Kerberos admin password - A string, the password for the administrator. This property is used only if creating service principals, and is not retained.

Changing services properties is documented in Setting Service Properties (BUI) and Setting Service Properties (CLI).

Kerberos Properties and Logs

The following table describes the mapping between Kerberos CLI properties and their BUI property descriptions.

> **Note:**
>
> Older Kerberos properties associated with the NFS service have been deprecated and will continue to function in scripts and workflows.

Table 3-18 Kerberos Properties

CLI Property	BUI Property
`realm`	Realm
`kdcs`	KDC(s)
`allow_weak_crypto`	Allow weak encryption types. Permits weak encryption types in Kerberos (arcfour-hmac-md5-exp, des-cbc-md5, and des-cbc-crc).
`principals`	Kerberos administrator principal
`principals - server`	Admin server - KDC host name
`principals - admin`	Admin principal - Administrator login name on KDC
`principals - password`	Password - Administrator login password on KDC
`importkeytab`	Import Keys - Imports keys in a keytab file from the KDC
`importkeytab - url`	URL of keytab file
`importkeytab - user`	User name for URL access
`importkeytab - password`	Password for URL access

The Kerberos service events log is available in `appliance-kit-kerberos:default`.

LDAP Configuration

Lightweight Directory Access Protocol (LDAP) is a directory service for centralizing management of users, groups, hostnames, and other resources (called objects). Oracle ZFS Storage Appliance can act as an LDAP client with the following characteristics:

- LDAP users can log in to the FTP and HTTP services.
- LDAP user names (instead of numerical IDs) can be used to configure root directory ACLs on a share.

  > **Note:**
  >
  > The appliance does not use LDAP authorization information. All authorization information is local.

- The LDAP server's certificate can be self-signed.
- You can supply a list of trusted CA certificates.

To configure LDAP, you configure schema settings, security settings, and LDAP servers as described in the following sections:

- LDAP Properties
- LDAP Custom Mappings

- Configuring LDAP Schema Settings - BUI, CLI
- LDAP Security Settings
- Configuring LDAP Security Settings - BUI, CLI
- Configuring the LDAP Server List - BUI, CLI
- Configuring LDAP Server Certificates
- Monitoring LDAP Server Status - BUI, CLI

If Active Directory (AD) is used for the appliance's LDAP service, you can automatically configure LDAP for the current AD domain by following the appropriate procedure: BUI or CLI.

After you have completed LDAP configuration, you can configure an existing LDAP user to be an appliance user. See Adding an LDAP User to the Appliance.

If AD is configured for the LDAP service and you are switching from domain to workgroup mode, you can optionally unconfigure or retain the LDAP configuration and the LDAP service. For more information, see Active Directory Domains and Workgroups.

LDAP Properties

For the appropriate settings for your environment, consult with your LDAP server administrator.

The tables in this section describe LDAP schema properties, security properties, and server properties.

Table 3-19   LDAP Schema Properties

BUI Property	CLI Property	Description
Base search DN	`base_dn`	The Distinguished Name of the base object, which is the starting point for directory searches. A default subtree specification is automatically prepended to this base search DN: `ou=people` for user searches, `ou=group` for group searches, `ou=netgroup` for netgroup searches. To override this default behavior, use the search descriptor properties listed below and described in LDAP Custom Mappings.
Search scope • One-level (non-recursive) • Subtree (recursive)	`search_scope` • `one` • `sub`	Which objects in the LDAP directory are searched, relative to the base object. For non-recursive, or `one`, search results are limited to only objects that are directly beneath the base search object. This is the default. For recursive, or `sub`, search results can include any object beneath the base search object.

Table 3-19 (Cont.) LDAP Schema Properties

BUI Property	CLI Property	Description
Schema definition for Users, Groups, and Netgroups • Search descriptor • Attribute mappings • Object class mappings	• `user_search`, `group_search`, `netgroup_search` • `user_mapattr`, `group_mapattr`, `netgroup_mapattr` • `user_mapobjclass`, `group_mapobjclass`, `netgroup_mapobjclass`	The schema used by the appliance. Use these properties to override the default search descriptor (base DN plus a default subtree specification), attribute mappings, and object class mappings for users, groups, and netgroups. For more information, see LDAP Custom Mappings.

Related Topics

- LDAP Custom Mappings
- Configuring LDAP Schema Settings - BUI, CLI

Table 3-20 LDAP Security Properties

BUI Property	CLI Property	Description
Authenticate as • Anonymous • Self • Proxy (Specific User)	`cred_level` • `anonymous` • `self` • `proxy`	Credentials used to authenticate the appliance to the LDAP server. See LDAP Security Settings for descriptions of these choices.
Enable SSL/TLS	`use_tls`	Toggles TLS (Transport Layer Security, the descendant of SSL) to establish secure connections to the LDAP server. If authenticating as Self, this option is not available because Self uses Kerberos encryption. If you specify port 636 when an LDAP server is added, the system configures LDAP and raw TLS. If you specify any other port when an LDAP server is added (typically 389), the system configures LDAP and StartTLS. When using raw TLS, a separate dedicated port is used for the secure TLS connection. With StartTLS, the LDAP server does not require a dedicated port to establish the encrypted LDAP connection; the LDAP server uses the same 389 port for a TLS connection.

Table 3-20 (Cont.) LDAP Security Properties

BUI Property	CLI Property	Description
Authentication Method - Simple (RFC 4513) - SASL/DIGEST-MD5	`auth_method` - `simple` - `sasl/DIGEST-MD5` - `sasl/GSSAPI` - `none`	Method used to authenticate the appliance to the LDAP server. If authenticating as Proxy, select the Simple or SASL/DIGEST-MD5 authentication method and set the DN and password. In the CLI, set `auth_method` to `sasl/GSSAPI` if authenticating as `self`. Set `auth_method` to `none` if authenticating as `anonymous`.
DN	`proxy_dn`	The distinguished name of the account that will be used for proxy authentication.
Password	`proxy_password`	The password for the proxy DN account.

Related Topics

- LDAP Security Settings
- Configuring LDAP Security Settings - BUI, CLI

Table 3-21 LDAP Server Properties

BUI Property	CLI Property	Description
- Use server order - Ignore server order	`use_server_order`	See the description of the server property for an explanation of the effect of the server order setting on a list of servers.
Server	`servers`	The list of LDAP servers to use. - If only one server is specified, the appliance uses only that server. If that server fails, LDAP services are unavailable. - If multiple servers are specified and Ignore server order is selected in the BUI or `use_server_order` is `false`, any functioning server on the list can be used at any time without preference. If any server fails, another server in the list is used. LDAP services remain available unless all specified servers fail. - If multiple servers are specified and Use server order is selected in the BUI or `use_server_order` is `true`, LDAP services will use the first available server on the list. The first server on the list is selected; if that server fails, the next server on the list is selected. LDAP services remain available unless all specified servers fail.

Related Topics

- Configuring LDAP Server Certificates
- Configuring the LDAP Server List - BUI, CLI
- Monitoring LDAP Server Status - BUI, CLI

LDAP Custom Mappings

To search the LDAP directory, the appliance uses a search descriptor that is the base search DN plus a prepended default subtree specification. The appliance also uses default object class names and default attribute names to find properties that are needed.

The appliance has the following LDAP search behavior, in accordance with RFC 2307:

- **User searches** – Prepends `ou=people` to the base search DN, uses object class `posixAccount`, and uses the attribute names shown in the first table.
- **Group searches** – Prepends `ou=group` to the base search DN, uses object class `posixGroup`, and uses the attribute names shown in the second table.
- **Netgroup searches** – Prepends `ou=netgroup` to the base search DN and uses object class `nisNetgroup`.

If these default values do not work with your environment, use the properties shown in the "LDAP Schema Properties" table in LDAP Properties to customize the search descriptor, object class names, and attribute names as shown in Configuring LDAP Schema Settings - BUI, CLI. To customize a search descriptor, enter the entire DN, including the base search DN and search scope. The appliance will use the customized value unmodified, and will ignore the values set for the base search DN and search scope properties. To customize object class names and attribute names, use *`default=new`* syntax, where *default* is the default value and *new* is the value that you want to use.

The following table shows the default attribute names that are used to find information about users.

Table 3-22 Attributes of the Users Data Type

Default Attribute Name	Description of Attribute Value
`uid`	User name. For example: `flastname`
`uidNumber`	Numeric user ID
`gidNumber`	Numeric primary group ID
`gecos`	Display name. For example: "*Firstname*, *Lastname*"

The following table shows the default attribute names that are used to find information about groups.

Table 3-23 Attributes of the Groups Data Type

Default Attribute Name	Description of Attribute Value
`cn`	Group name
`gidNumber`	Numeric group ID
`memberUid`	List of usernames of members

Related Topics

- "LDAP Schema Properties" table in LDAP Properties
- Configuring LDAP Schema Settings - BUI, CLI

Configuring LDAP Schema Settings (BUI)

Use the following procedure to configure LDAP schema settings.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **LDAP**.

   On the **Properties** tab, go to the **Schema** section of the page.
3. Enter a value for **Base search DN**.

   For example, enter:

   ```
   dc=example,dc=com
   ```

   This base search DN is automatically prepended with `ou=people` for user searches, `ou=group` for group searches, and `ou=netgroup` for netgroup searches. If these values do not work in your environment, set different values for search descriptor in the schema definition for that type of search. The value of base search DN is ignored if you provide a value for search descriptor in the schema definition.
4. Select either **recursive** or **non-recursive** search scope.

   This selection applies to all searches. To override this selection for specific types of searches, provide a value for search descriptor in the schema definition for that type of search. The value of search scope is ignored if you provide a value for search descriptor in the schema definition.
5. Click **Edit** next to the schema definition heading.

   The **Edit LDAP Schema Definition** dialog box opens. The dialog box has three tabs: **Users**, **Groups**, and **Netgroups**. Each tab has three property fields: **Search descriptor**, **Attribute mappings**, and **Object class mappings**.
6. Edit the **Search descriptor** fields.

   The default search descriptor that is used for user searches is `ou=people,`*`base-search-DN`* . The default search descriptor that is used for group searches is `ou=group,`*`base-search-DN`* . The default search descriptor that is used for netgroup searches is `ou=netgroup,`*`base-search-DN`* . If your LDAP database does not have subtrees named `people`, `group`, or `netgroup`, then searches of the database will fail as object not found.

   Edit the search descriptor fields on each tab to enter the correct subtrees to search for users, groups, and netgroups. For example, on the **Users** tab, you might enter the following for the search descriptor:

   ```
   ou=employees,dc=example,dc=com
   ```

   If your LDAP database does not have subtrees for users and groups, use this search descriptor field to re-enter the base search DN to prevent `ou=people` or `ou=group` from being prepended automatically. For example, enter the following in the **Users** or **Groups** search descriptor field:

   ```
   dc=example,dc=com
   ```

   You must include the full base search DN in the search descriptor value. Also include your scope selection. Both the base search DN and scope selection will be ignored and the search descriptor value will be used instead. The example in the previous paragraph specifies non-recursive search. To specify recursive search, change that example to the following:

```
dc=example,dc=com?sub
```

7. Edit the **Attribute mappings** fields.

   The default attributes that are used for user searches are shown in table "Attributes of the Users Data Type" in LDAP Custom Mappings. The default attributes that are used for group searches are shown in table "Attributes of the Groups Data Type" in LDAP Custom Mappings.

   If your organization stores this data in different attributes, use the **Attribute mappings** field to specify the attribute to use to retrieve the given data. For example, to use `employeename` instead of `uid` as the attribute for user names, enter `uid=employeename` in the attribute mappings field on the Users tab.

   To specify additional attribute changes for a given data type, click the add icon ⊕ to the right of the **Attribute mappings** field.

8. Edit the **Object class mappings** fields.

   The default object class that is used for user searches is `posixAccount`. The default object class that is used for group searches is `posixGroup`. The default object class that is used for netgroup searches is `nisNetgroup`.

   If your environment uses a different object class, use the **Object class mappings** field to specify the name of the object class to use. For example, to use `unixaccount` instead of `posixAccount` as the user object class, enter `posixAccount=unixaccount` in the **Object class mappings** field on the **Users** tab.

   To specify additional object class changes, click the add icon ⊕ to the right of the **Object class mappings** field.

9. Click **Save** in the **Edit LDAP Schema Definition** dialog box.

10. Click **APPLY** at the top of the LDAP page.

Related Topics

- "LDAP Schema Properties" table in LDAP Properties
- "Attributes of the Users Data Type" table in LDAP Custom Mappings
- "Attributes of the Groups Data Type" table in LDAP Custom Mappings

Configuring LDAP Schema Settings (CLI)

Use the following procedure to configure LDAP schema settings.

1. Go to `configuration services ldap` and enter the `show` command.

2. Enter a value for the base search DN.

   Use quotation marks around the value to preserve the embedded equal symbols and comma. For example, enter:

   ```
   hostname:configuration services ldap> set base_dn="dc=example,dc=com"
   ```

   This base search DN is automatically prepended with `ou=people` for user searches, `ou=group` for group searches, and `ou=netgroup` for netgroup searches. If these values do not work in your environment, set different values for `user_search`, `group_search`, or netgroup_search. The value of base_dn is ignored if you provide a value for user_search, group_search, or netgroup_search.

3. Specify either recursive or non-recursive scope.

   For recursive search, set search_scope to sub. For non-recursive search, set search_scope to one.

   The value of search_scope applies to all searches. To override search_scope for specific types of searches, provide a value for user_search, group_search, or netgroup_search.

4. Set search descriptor properties.

   The search descriptor properties are user_search, group_search, and netgroup_search.

   By default, ou=people is prepended to the value of base_dn for user searches, ou=group is prepended to the value of base_dn for group searches, and ou=netgroup is prepended to the value of base_dn for netgroup searches. If your LDAP database does not have subtrees named people, group, or netgroup, then searches of the database will fail as object not found.

   Set the values of the search descriptor properties to specify the correct subtrees to search for users, groups, and netgroups. Use quotation marks around the value to preserve the embedded equal symbols and commas. For example, you might enter the following for the search descriptor for users:

   ```
   hostname:configuration services ldap> set
   user_search="ou=employees,dc=example,dc=com"
   ```

   If your LDAP database does not have subtrees for users and groups, use the search descriptor properties to re-enter the base search DN to prevent ou=people or ou=group from being prepended automatically, as shown in the following example:

   ```
   hostname:configuration services ldap> set group_search="dc=example,dc=com"
   ```

   You must include the value of base_dn in the search descriptor value. Also include your scope selection. Both base_dn and search_scope will be ignored and the search descriptor value will be used instead. The example in the previous paragraph specifies non-recursive search. To specify recursive search, change that example to the following:

   ```
   hostname:configuration services ldap> set group_search="dc=example,dc=com?sub"
   ```

5. Set the attribute mapping properties.

   The default attributes that are used for user searches are shown in table "Attributes of the Users Data Type" in LDAP Custom Mappings. The default attributes that are used for group searches are shown in table "Attributes of the Groups Data Type" in LDAP Custom Mappings.

   If your environment stores this data in different attributes, use the user_mapattr, group_mapattr, and netgroup_mapattr properties to specify the attribute to use to retrieve the given data. For example, enter the following command to use employeename instead of uid as the attribute for user names:

   ```
   hostname:configuration services ldap> set user_mapattr="uid=employeename"
   ```

   Enclose the mapping in quotation marks to preserve the equal symbol in the mapping.

   To specify multiple attribute changes, enter a list of mappings with quotation marks around each mapping, as shown in the following example:

   ```
   hostname:configuration services ldap> set
   user_mapattr="uid=employeename","uidNumber=employeenumber"
   ```

6. Set the object mapping properties.

   The default object class that is used for user searches is `posixAccount`. The default object class that is used for group searches is `posixGroup`. The default object class that is used for netgroup searches is `nisNetgroup`.

   If your environment uses a different object class, use the `user_mapobjclass`, `group_mapobjclass`, and `netgroup_mapobjclass` properties to specify the name of the object class to use. For example, enter the following command to use `unixaccount` instead of `posixAccount` as the user object class:

   ```
   hostname:configuration services ldap> set
   user_mapobjclass="posixAccount=unixaccount"
   ```

7. Commit changed property settings.

   Enter `show`, review property settings, and enter `commit`.

Related Topics

- "LDAP Schema Properties" table in LDAP Properties
- "Attributes of the Users Data Type" table in LDAP Custom Mappings
- "Attributes of the Groups Data Type" table in LDAP Custom Mappings

LDAP Security Settings

> **Caution:**
>
> To reduce security risks, always configure LDAP with SSL/TLS or Kerberos.

To configure security settings for the LDAP service, first specify the credentials to use to authenticate the appliance to the LDAP server. Then specify other properties as necessary to support the credentials choice.

The appliance can authenticate by using one of the following sets of credentials:

- **Anonymous:**
  - Anonymous authentication restricts data access for the appliance to only data that is available to everyone.
  - You can choose to enable the TLS (formerly known as SSL) protocol. Enabling TLS is highly recommended so that critical information is sent securely.
- **Self:** Self authentication uses the user's identity and credentials to authenticate the appliance. Self authentication uses Kerberos encryption and the SASL/GSSAPI authentication method.
- **Proxy (Specific User):**
  - Proxy authentication uses a proxy for a specific user account.
  - You can choose to enable the TLS (formerly known as SSL) protocol. Enabling TLS is highly recommended so that critical information is sent securely.
  - You must select the authentication method: either Simple (RFC 4513) or SASL/DIGEST-MD5.

   - You must specify the proxy DN and the proxy password. The proxy DN is the distinguished name of the account that will be used for proxy authentication. The proxy password is the password for the proxy DN account.

If you specify port 636 when an LDAP server is added, the system configures LDAP and raw TLS. If you specify any other port when an LDAP server is added (typically 389), the system configures LDAP and StartTLS. For information about StartTLS, see "LDAP Security Properties" table in LDAP Properties.

Configuring LDAP Security Settings (BUI)

Use the following procedure to configure LDAP security settings.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **LDAP**.

   On the **Properties** tab, go to the **Security Settings** section of the page.
3. Specify the credentials to use to authenticate the appliance to the LDAP server.

   For **Authenticate as**, select **Anonymous**, **Self**, or **Proxy (Specific User)**.
4. Enable SSL/TLS.

   If you selected either **Anonymous** or **Proxy** for **Authenticate as**, you can choose to enable SSL/TLS. Enabling TLS is highly recommended so that critical information is sent securely.
5. Specify an **Authentication Method**.

   If you selected **Proxy** for **Authenticate as**, select either **Simple (RFC 4513)** or **SASL/DIGEST-MD5** from the **Authentication Method** menu. Then specify the **DN** and **Password**.
6. Click **APPLY**.

   The LDAP server configuration is validated. If the `Proxy DN` or `Proxy Password` fails or times out, a warning is displayed.

Related Topics

- "LDAP Security Properties" table in LDAP Properties
- LDAP Security Settings

Configuring LDAP Security Settings (CLI)

Use the following procedure to configure LDAP security settings.

1. Go to `configuration services ldap` and enter `show` to view the properties.

   The following table shows property value combinations that are valid for the remaining steps in this procedure.

cred_level	auth_method	use_tls
anonymous	none	true
anonymous	none	false
self	sasl/GSSAPI	false
proxy	simple	true

cred_level	auth_method	use_tls
proxy	simple	false **Note:** This setting is permitted, but not recommended because the user's distinguished name (DN) and password will be sent in plain text.
proxy	sasl/DIGEST-MD5	true
proxy	sasl/DIGEST-MD5	false

2. Specify the credentials to use to authenticate Oracle ZFS Storage Appliance to the LDAP server.

   Set `cred_level` to `anonymous`, `self`, or `proxy`.

   ```
   hostname:configuration services ldap> set cred_level=proxy
   ```

3. Specify an authentication method.

   Set `auth_method` to one of the following options:

   - `none` - None (use with `anonymous`)
   - `sasl/GSSAPI` - SASL/GSSAPI (use with `self`)
   - `simple` - Simple, RFC 4513 (use with `proxy`)
   - `sasl/DIGEST-MD5` - SASL/DIGEST-MD5 (use with `proxy`)

   ```
   hostname:configuration services ldap> set auth_method=sasl/DIGEST-MD5
   ```

4. Set additional properties for proxy credentials.

   If `cred_level` is set to `proxy`, then set the proxy account name and password.

   ```
   hostname:configuration services ldap> set proxy_dn=ProxyName
   hostname:configuration services ldap> set proxy_password=MyPassword5
   ```

5. Enable SSL/TLS.

   If you specified either `anonymous` or `proxy` for `cred_level`, you can choose to enable SSL/TLS. Enabling TLS is highly recommended so that critical information is sent securely.

   ```
   hostname:configuration services ldap> set use_tls=true
   ```

6. Enter `commit`.

   Changes to the LDAP server configuration will be validated when committed. If the `proxy_dn` or `proxy_password` validation fails or times out, a warning message is displayed.

Related Topics

- "LDAP Security Properties" table in LDAP Properties
- LDAP Security Settings

Configuring the LDAP Server List (BUI)

Use the following procedure to configure the LDAP server list.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **LDAP**.

   On the **Properties** tab, scroll to the **LDAP Servers** section of the page.
3. Configure the list.

- If you plan to have more than one server in the list, decide whether you want the servers to be tried in a particular order.

  If the servers should be tried in the order in which they are listed, select the **Use server order** button at the top of the list. By default, **Ignore server order** is selected. For more information about how the servers are used, see table "LDAP Server Properties" in LDAP Properties.

- To add a server to the list, click the add icon ⊕ to the left of the **LDAP Servers** section title.

  In the **New LDAP Server** dialog box, enter the server name, such as *`hostname`*`.example.com`, and optionally the port.

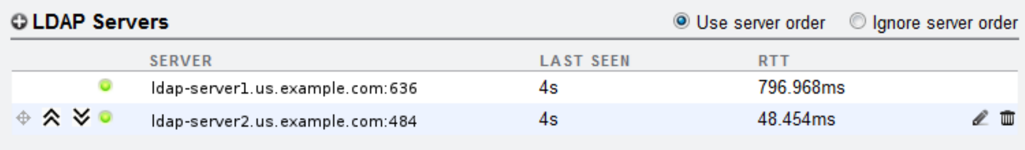

⊕ LDAP Servers ● Use server order ○ Ignore server order

	SERVER	LAST SEEN	RTT	
●	ldap-server1.us.example.com:636	4s	796.968ms	
✥ ⇑ ⇓ ●	ldap-server2.us.example.com:484	4s	48.454ms	✎ 🗑

- To change the order of the servers in the list, use the tools on the left of the row. These tools are only available when **Use server order** is selected and applied. Click on the row of the server that you want to move and do one of the following:
  - Click the up or down arrows to the left of the row.
  - Click and hold on the move icon to the far left of the row and drag the server entry up or down the list.
- To change the name of a server on the list, click on the row of the server that you want to modify, and click the edit icon ✎ to the right of the row.
- To remove a server from the list, click on the row of the server that you want to remove, and click the trash icon 🗑 to the right of the row.

4. Click **APPLY** at the top of the page.

Related Topics

- "LDAP Server Properties" table in LDAP Properties
- Configuring LDAP Server Certificates
- Monitoring LDAP Server Status (BUI)

Configuring the LDAP Server List (CLI)

Use the following procedure to configure the LDAP server list.

1. Go to `configuration services ldap`.
2. Enter the `list` command to show existing LDAP servers.
3. Configure the list.
   - If you plan to have more than one server in the list, decide whether you want the servers to be tried in a particular order. If the servers should be tried in the order in which they are specified in the `servers` property, set the `use_server_order` property to `true`. By default, the value of the `use_server_order` property is `false`. For more information about how the servers are used, see table "LDAP Server Properties" in LDAP Properties.

- Check the value of the `servers` property.

  ```
  hostname:configuration services ldap> get servers
            servers = ldap-server2.us.example.com:484,ldap-
  server1.us.example.com:636
  ```

- To add, remove, or reorder servers in the list, reset the value of the `servers` property. The following example reorders the list:

  ```
  hostname:configuration services ldap> set servers=ldap-
  server1.us.example.com:636,ldap-server2.us.example.com:484
  ```

4. Enter `commit`.

Related Topics

- "LDAP Server Properties" table in LDAP Properties
- Configuring LDAP Server Certificates
- Monitoring LDAP Server Status (CLI)

Configuring LDAP Server Certificates

An LDAP server's certificate can be CA-signed or self-signed. This section describes how to initially configure certificates and how to manage a new certificate when the previous certificate expires.

Initially Configuring LDAP Server Certificates

For more information about trusted certificates, see the sections about trusted certificates in Configuring Certificates.

You can supply a list of trusted CA certificates. LDAP server certificates issued by those trusted CAs do not require special management.

If an LDAP server's certificate is not issued by a trusted CA, whether the certificate is issued by a CA or is self-signed, you will be asked to review and approve the certificate. The **Accept LDAP Server Certificate** dialog box displays information about the certificate and requests that you accept or reject the certificate. If you accept the certificate, that certificate is added to the list of trusted certificates.

Managing Expired and New LDAP Server Certificates

If you individually accepted a certificate, either a CA-signed certificate or a self-signed certificate, then when the LDAP server's certificate expires, you must approve the new certificate. Select the server, test the connection, and examine and approve the new certificate. See Approving a New LDAP Server Certificate - BUI, CLI.

If you supply CA certificates, changes in the individual server certificates are handled automatically. When your server changes CA certificates, ensure that the new CA certificate is added to the appliance before your LDAP server starts using it. If the server starts using the new CA certificate before you add it to the appliance, your LDAP service will be interrupted.

Approving a New LDAP Server Certificate (BUI)

Use this procedure to accept a new certificate after the previous certificate expired.

Before You Begin

TLS must be enabled. Make sure the Enable SSL/TLS box is checked.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **LDAP**.

   On the **Properties** tab, scroll to the **LDAP Servers** section of the page.
3. In the table, click the edit icon ✎ for the server that has the new certificate.
4. Click the **Test Connection** button in the **Edit LDAP Server** dialog box to test the TLS connection.

   A new dialog reports whether the new certificate is trusted.
5. Click **OK** in the trusted certificate dialog box.

   If the trusted certificate dialog box reported the certificate is not trusted, the **Accept LDAP Server Certificate** dialog box opens. This dialog box displays information about the certificate, and has **Reject** and **Accept** buttons.
6. Review the certificate information, and click **Accept**.

   The certificate is added to the list of trusted certificates.

Related Topics

Viewing Trusted Certificate Details (BUI)

Approving a New LDAP Server Certificate (CLI)

Use this procedure to accept a new certificate after the previous certificate expired.

Before You Begin

TLS must be enabled. Make sure the `use_tls` property is set to `true`.

1. Go to `configuration services ldap`.
2. Enter the `list` command to show the list of LDAP servers.
3. Select a server.
4. Enter the `test` command to test the TLS connection.

   Information about the new certificate is displayed.
5. Examine and approve the new certificate.

   The certificate is added to the list of trusted certificates. If you enter the `test` command again, the message `Certificate is trusted` is displayed.

Related Topics

Viewing Trusted Certificate Details (CLI)

Monitoring LDAP Server Status (BUI)

Use the following procedure to monitor the LDAP server status.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **LDAP**.
3. View server status.

   On the **Properties** tab, scroll to the **LDAP Servers** section of the page.

   For each server, the following status is displayed:

- **Status icon** - Indicator icon that represents the status of the server. The indicator icon is either online or unavailable ●.
- **Last Seen** - The total time since the last response was received from each LDAP server.
- **RTT** - The round-trip time to get the response from the server.

4. View the LDAP server logs.

   a. Click the **Logs** tab at the top of the **LDAP** page.

   b. Select an LDAP server from the drop-down menu.

      The log entries contain a time and description for specific LDAP server alerts.

   c. View additional logs.

      Use the arrow buttons to the right of the server menu to view previous or later logs for the selected server.

Related Topics

- "LDAP Server Properties" table in LDAP Properties
- Configuring the LDAP Server List (BUI)

Monitoring LDAP Server Status (CLI)

Use the following procedure to monitor the LDAP server status.

1. Go to `configuration services ldap`.
2. Enter the `list` command to show the status of LDAP servers.

```
hostname:configuration services ldap> list
SERVER       STATUS       LDAP SERVER
server-000   unavailable  ldap-server1.us.example.com:636
server-001   online       ldap-server2.us.example.com:484
```

Related Topics

- "LDAP Server Properties" table in LDAP Properties
- Configuring the LDAP Server List (CLI)

Adding an LDAP User to the Appliance

After you have completed LDAP configuration, you can configure existing LDAP users to be able to log in to Oracle ZFS Storage Appliance using their LDAP username and password.

> **Note:**
>
> The existing LDAP user ID cannot be less than 100, cannot be greater than 2147483646, and cannot be equal to 60001, 60002, or 65534. Those UIDs are reserved by the operating system vendor for use in future applications. Their use by end system users or vendors of layered products is not supported and can cause security issues with other applications.

In the BUI, follow procedure Adding an Administrator or User (BUI), using the following property values:

1. Select **Directory** from the **Type** drop-down menu.
2. For **Username**, specify the user's existing LDAP username.

In the CLI, follow procedure Adding an Administrator or User (CLI), using the `directory` user type and the user's existing LDAP username.

```
hostname:configuration users> directory LDAPusername
```

User ID (`uid`) and Password (`initial_password`) are set automatically: LDAP values are used, or if NIS is configured, the NIS values are used.

> **Note:**
>
> If both NIS and LDAP are configured on the appliance and the services return different information for a particular user, the appliance uses the data provided by NIS.

NDMP Configuration

The Network Data Management Protocol (NDMP) service enables the system to participate in NDMP-based backup and restore operations controlled by a remote NDMP client called a Data Management Application (DMA). Using NDMP, data stored in administrator-created shares on the appliance can be backed up and restored to both locally attached tape devices and remote systems. Locally-attached tape devices can also be exposed to the DMA for backing up and restoring remote systems.

NDMP cannot be used to back up and restore system configuration data. Instead, use the Configuration Backup and Restore feature (see Backing Up the Configuration in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*).

To configure NDMP, see the following sections:

- NDMP Local vs. Remote Configurations
- NDMP Backup Formats and Types
- NDMP Backup with Types dump and tar
- NDMP Backup with Type zfs
- NDMP Incremental Backups
- Replica Backups
- NDMP Properties and Logs

NDMP Local vs. Remote Configurations

Oracle ZFS Storage Appliance supports backup and restore using both a *local* configuration, in which tape drives are physically attached to the appliance, and a *remote* configuration, in which data is streamed to another system on the same network. In both cases, the backup must be managed by a supported DMA.

In local configurations, supported tape devices, including both drives and changers (robots), are physically connected to the system using a supported SCSI or Fibre Channel (FC) card configured in Initiator mode. These devices can be viewed on the NDMP Status screen. The NDMP service presents these devices to a DMA when the DMA scans for devices. Once configured in the DMA, these devices are available for backup and restore of the appliance or other systems on the same network. After adding tape drives or changers to the system or removing such devices from the system, a reboot may be required before the changes will be recognized by the NDMP service. After that, the DMA may need to be reconfigured because tape device names may have changed.

In remote configurations, the tape devices are not physically connected to the system being backed up and restored (the data server) but rather to the system running the DMA or a separate system (the tape server). These are commonly called "3-way configurations" because the DMA controls two other systems. In these configurations, the data stream is transmitted between the data server and the tape server over an IP network.

NDMP Backup Formats and Types

The NDMP protocol does not specify a backup data format. The appliance supports three backup types corresponding to different implementations and on-tape formats. DMAs can select a backup type using the following values for the NDMP environment variable `TYPE`:

Table 3-24 NDMP Backup Formats and Types

Backup Type	Details
dump	File-based for filesystems only. Supports file history and direct access recovery (DAR).
tar	File-based for filesystems only. Supports file history and direct access recovery (DAR).
zfs	Share-based for both filesystems and volumes. Does not support file history or direct access recovery (DAR), but may be faster for some datasets. Only supported with NDMPv4.

There is no standard NDMP data stream format, so backup streams generated on the appliance can only be restored on Oracle ZFS Storage Appliance systems running compatible software. Future versions of appliance software can generally restore streams backed up from older versions of the software, but the reverse is not necessarily true. For example, the "zfs" backup type was new in 2010.Q3, and systems running 2010.Q1 or earlier cannot restore backup streams created using type "zfs" under 2010.Q3.

NDMP Backup with Types dump and tar

When backing up with "dump" and "tar" backup types, administrators specify the data to backup by a filesystem path, called the *backup path*. For example, if the administrator configures a backup of `/export/home`, then the share mounted at that path will be backed up. Similarly, if a backup stream is restored to `/export/code`, then that's the path where files will be restored, even if they were backed up from another path.

Only paths that are mountpoints of existing shares, or contained within existing shares, may be specified for backup. If the backup path matches a share's mountpoint, only that share is backed up. Otherwise the path must be contained within a share, in which case only the portion of that share under that path is backed up. In both cases, other shares mounted inside the specified share under the backup path will not be backed up; these shares must be specified separately for backup.

Snapshots - If the backup path specifies a live filesystem (such as `/export/code`) or a path contained within a live filesystem (such as `/export/code/src`), the appliance immediately takes a new snapshot and backs up the given path from that snapshot. When the backup completes, the snapshot is destroyed. If the backup path specifies a snapshot (for example, /

`export/code/.zfs/snapshot/mysnap`), no new snapshot is created and the system backs up from the specified snapshot.

Share metadata - To simplify backup and restore of complex share configurations, "dump" and "tar" backups include share metadata for projects and shares associated with the backup path. This metadata describes the share configuration on the appliance, including protocol sharing properties, quota properties, and other properties configured on the Shares screen. This is not to be confused with filesystem metadata like directory structure and file permissions, which is also backed up and restored with NDMP.

For example, if you back up `/export/proj`, the share metadata for all shares whose mountpoints start with /export/proj will be backed up, as well as the share metadata for their parent projects. Similarly, if you back up `/export/someshare/somedir`, and a share is mounted at `/export/someshare`, that share and its project's share metadata will be backed up.

When restoring, if the destination of the restore path is not contained inside an existing share, projects and shares in the backup stream will be recreated as needed with their original properties as stored in the backup. For example, if you back up /export/foo, which contains project `proj1` and shares `share1` and `share2`, and then destroy the project and restore from the backup, then these two shares and the project will be recreated with their backed-up properties as part of the restore operation.

During a restore, if a project exists that would have been automatically recreated, the existing project is used and no new project is automatically created. If a share exists that would have been automatically recreated, and if its mountpoint matches what the appliance expects based on the original backup path and the destination of the restore, then the existing share is used and no new share is automatically created. Otherwise, a new share is automatically created from the metadata in the backup. If a share with the same name already exists (but has a different mountpoint), then the newly created share will be given a unique name starting with `ndmp-` and with the correct mountpoint.

It is recommended that you either restore a stream whose datasets no longer exist on the appliance, allowing the appliance to recreate datasets as specified in the backup stream, or precreate a destination share for restores. Either of these practices avoids surprising results related to the automatic share creation described earlier.

NDMP Backup with Type zfs

When backing up with type "zfs", administrators specify the data to backup by its canonical name on Oracle ZFS Storage Appliance. This can be found underneath the name of the share in the BUI:

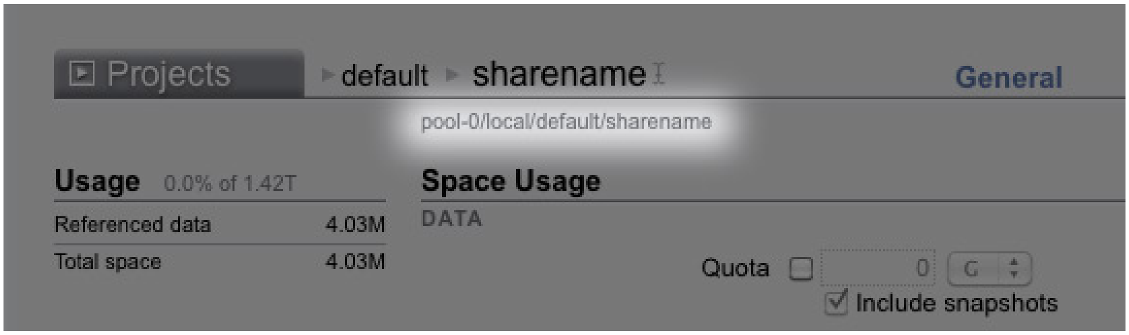

pool-0/local/default/sharename

Usage 0.0% of 1.42T

Referenced data	4.03M
Total space	4.03M

Quota □ 0 G

☑ Include snapshots

or in the CLI as the value of the `canonical_name` property. Canonical names do not begin with a leading '/', but when configuring the backup path, the canonical name must be prefixed with '/'.

Both projects and shares can be specified for backup using type "zfs". If the canonical name is specified as-is, then a new snapshot is created and used for the backup. A specific snapshot can be specified for backup using the `@snapshot` suffix, in which case no new snapshot is created and the specified snapshot is backed up. For example:

Table 3-25 Canonical Names and Shares Backed Up

Canonical Name	Shares Backed Up
`pool-0/local/default`	New snapshot of the local project called `default` and all of its shares.
`pool-0/local/default@yesterday`	Named snapshot `yesterday` of local project `default`, and all of its shares having snapshot `yesterday`.
`pool-0/local/default/code`	New snapshot of share `code` in local project `default`. `code` could be a filesystem or volume.
`pool-0/local/default/code@yesterday`	Named snapshot `yesterday` of share `code` in local project `default`. `code` could be a filesystem or volume.

Because level-based incremental backups using the "zfs" backup type require a base snapshot from the previous incremental, the default behavior for level backups for which a new snapshot is created is to keep the new snapshot so that it can be used for subsequent incremental backups. If the DMA indicates that the backup will not be used for subsequent incremental backups by setting `UPDATE=n`, the newly created snapshot is destroyed after the backup. Existing user snapshots are never destroyed after a backup. For details, see NDMP Incremental Backups.

Share metadata - Share metadata (that is, share configuration) is always included in "zfs" backups. When restoring a full backup with type "zfs", the destination project or share must not already exist. It will be recreated from the metadata in the backup stream. When restoring an incremental backup with type "zfs", the destination project or share must already exist. Its properties will be updated from the metadata in the backup stream. For details, see NDMP Incremental Backups.

NDMP Incremental Backups

Oracle ZFS Storage Appliance supports level-based incremental backups for all of the above backup types. To specify a level backup, DMAs typically specify the following three environment variables:

Variable	Details
`LEVEL`	Integer from 0 to 9 identifying the backup level.
`DMP_NAME`	Specifies a particular incremental backup set. Multiple sets of level incremental backups can be used concurrently by specifying different values for `DMP_NAME`.
`UPDATE`	Indicates whether this backup can be used as the base for subsequent incremental backups

By definition, a level-N backup includes all files changed since the previous backup of the same backup set (specified by `DMP_NAME`) of the same share using `LEVEL` less than *N*. `Level-0` backups always include all files. If UPDATE has value y (the default), then the current backup is recorded so that future backups of level greater than *N* will use this backup as a base. These variables are typically managed by the DMA, and do not need to be configured directly by administrators. The following table describes a sample incremental backup schedule:

Table 3-26 Sample Incremental Backup Schedule

Day	Details
First of month	Level-0 backup. Backup contains all files in the share.
Every 7th, 14th, 21st of month	Level-1 backup. Backup contains all files changed since the last full (monthly) backup
Every day	Level-2 backup. Backup contains all files changed since the last level-1 backup

To recover the filesystem's state as it was on the 24th of the month, an administrator typically restores the Level-0 backup from the 1st of the month to a new share, then restores the Level-1 backup from the 21st of the month, and then restores the Level-2 backup from the 24th of the month.

To implement level-based incremental backups the appliance must keep track of the level backup history for each share. For "tar" and "dump" backups, the level backup history is maintained in the share metadata. Incremental backups traverse the filesystem and include files modified since the time of the previous level backup. At restore time, the system simply restores all the files in the backup stream. In the above example, it would therefore be possible to restore the Level-2 backup from the 24th onto any filesystem and the files contained in that backup stream will be restored even though the target filesystem may not match the filesystem where the files were backed up. However, best practice suggests using a procedure like the above which starts from an empty tree restores the previous level backups in order to recover the original filesystem state.

To implement efficient level-based incremental backups for type "zfs", the system uses a different approach. Backups that are part of an incremental set do not destroy the snapshot used for the backup but rather leave it on the system. Subsequent incremental backups use this snapshot as a base to quickly identify the changed filesystem blocks and generate the backup stream. As a consequence, the snapshots left by the NDMP service after a backup must not be destroyed if you want to create subsequent incremental backups.

Another important consequence of this behavior is that in order to restore an incremental stream, the filesystem state must exactly match its state at the base snapshot of the incremental stream. In other words, in order to restore a Level-2 backup, the filesystem must look exactly as it did when the previous Level-1 backup completed. Note that the above commonly used procedure guarantees this because when restoring the Level-2 backup stream from the 24th, the system is exactly as it was when the Level-1 backup from the 21st completed because that backup has just been restored.

The NDMP service will report an error if you attempt to restore an incremental "zfs" backup stream to a filesystem whose most recent snapshot doesn't match the base snapshot for the incremental stream, or if the filesystem has been changed since that snapshot. You can configure the NDMP service to rollback to the base snapshot just before the restore begins by specifying the NDMP environment variable ZFS_FORCE with value y or by configuring the Rollback datasets property of the NDMP service (see NDMP Properties and Logs).

Replica Backups

Oracle ZFS Storage Appliance supports direct backup of replicas and replica snapshots with the "zfs" backup type. It is not necessary to first clone a replica dataset (project or share) in order to back it up.

> **Note:**
>
> Because the backup is of a replica, the source dataset properties are backed up rather than those of the target.

Enabling Replica Backups

To enable replica backups, apply the corresponding deferred update. For more information, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

Replica backups require software version 2011.1.0 or later on the source. If the replica backup will be restored to the source with the original replicated dataset, the source must run software version 2013.1.4 or later.

Replica Backup Syntax

To back up a replicated project or share, input the ZFS dataset name without a snapshot extension into the DMA. `ndmpd` uses the Oracle ZFS Storage Appliance software to determine the latest complete replica snapshot to back up. To specify a replica dataset for backup, use copy and paste to avoid mistyping long replica dataset names, which may include a UUID.

If a user-generated snapshot extension is included, `ndmpd` backs up the indicated user snapshot. If a system-generated extension is included (begins with `.rr`), the backup fails and generates a message that is logged to the DMA console.

Replica Backup Persistent Holds

Persistent holds are taken on backed-up snapshots when the backups complete. This is necessary for future incremental backups, which use current snapshots as a base, otherwise the replication subsystem may delete replica snapshots it no longer needs. Holds are released by `ndmpd` when the snapshots are no longer needed.

Persistent holds can be cleared manually. When deleting a replica snapshot with a hold on it, a confirmation is displayed warning of the potential impact to ongoing or future NDMP backups. Snapshots required by the replication subsystem cannot be deleted.

If incremental backups are not needed, prevent persistent holds by setting the DMA `UPDATE` parameter to no `(UPDATE=n)`. `UPDATE=y` is the default mode. For more information about the `UPDATE` NDMP environment variable, see the technical brief NDMP Implementation Guide for the Oracle ZFS Storage Appliance (https://www.oracle.com/storage/technologies/nas-unified-storage-documentation.html).

Incremental Replica Backups

Continuing (incrementing) a backup series across a replication reversal or sever is not supported; instead, start a new backup series. Use a full (`Level-0`) backup for the first backup after a replication state change has occurred, such as on a new source after a reversal or sever has taken place.

Static snapshot extensions that do not change per level are NOT supported for user-generated replica snapshots (snapshots not starting with `.rr`). This prevents name collisions, which generate an error and can cause replication to fail.

Some DMAs do not support zfs-type replica incremental backup and restore operations for snapshot extension name changes per level. To conserve appliance space and ensure that such snapshots are not preserved for future incremental backups, set `UPDATE=n` at the time of backup of replica. User-generated snapshots can be removed manually.

Even if no user data has changed in a restored dataset, changed metadata can cause incremental replica restores to fail. To avoid this, always roll back to the base snapshot before incremental replica restores by setting the `ZFS rollback before restore` parameter to `Always`.

For non-incremental replica backups, such as for one-off backups, set `UPDATE=n` so future snapshots are not saved and consume space. Some older replica snapshots preserved for future incremental backups, such as those created by setting `UPDATE=y`, may no longer be needed and waste space. These snapshots are safe to manually destroy. Snapshots needed by the replication subsystem cannot be deleted. Unneeded snapshots can be deleted after confirming the warning message about possible impacts to ongoing or future NDMP backups if the snapshot is deleted.

NDMP Properties and Logs

The NDMP service configuration consists of the following properties and logs.

Table 3-27 NDMP Properties

Property	Description
Version	The version of NDMP that your DMA supports.
TCP port (v4 only)	The NDMP default connection port is 10000. NDMPv3 always uses this port. NDMPv4 allows a different port if needed.
Ignore metadata-only changes	Directs the system to backup only files in which content has changed, ignoring files for which only metadata, such as permissions or ownership, has changed. This option only applies to incremental "tar" and "dump" backups and is disabled by default.
Target restore pool(s)	When you perform a full restore using "tar" or "dump", the system re-creates datasets if there is no share mounted at the target. Because the NDMP protocol specifies only the mountpoint, the system chooses a pool in which to recreate projects and shares. On a system with multiple pools, this property lets you specify one or more pools. Multiple pools only need to be specified in a cluster with active pools on each head. You must ensure that this list is kept in sync with any storage configuration changes. If none of the pools exist or are online, the system will select a default pool at random.
Allow token-based backup	Enables or disables token-based method for ZFS backup. This property is off by default.

Table 3-27 (Cont.) NDMP Properties

Property	Description
ZFS rollback before restore (v4 only)	Only applies to backups with type "zfs". Determines whether when restoring an incremental backup the system rolls back the target project and share to the snapshot used as the base for the incremental restore. If the project and shares are rolled back, then any changes made since that snapshot will be lost. This setting is normally controlled by the DMA via the `ZFS_FORCE` environment variable, but this property can be used to override the DMA setting to always rollback these data sets or never roll them back. For details, see NDMP Incremental Backups. Not rolling them back will cause the restore to fail unless they have already been manually rolled back. This property is intended for use with DMAs that do not allow administrators to configure custom environment variables like `ZFS_FORCE`.
Allow direct access recovery	Enables the system to locate files by position rather than by sequential search during restore operations. Enabling this option reduces the time it takes to recover a small number of files from many tapes. You must specify this option at backup time in order to be able to recover individual files later.
Restore absolute paths (v3 only)	Specifies that when a file is restored, the complete absolute path to that file is also restored (instead of just the file itself). This option is disabled by default.
Share creation on restore	Configures the restore operation to create a new share based on the backup type: • **All** - Allows all backup types to create a new share on restore • **Tar_Dump** - Allows backup types "tar" and "dump" to create a new share on restore • **ZFS** - Allows backup type "zfs" to create a new share on restore • **None** - No backup type can create a new share on restore
DMA tape mode (for locally attached drives)	Specifies whether the DMA expects System V or BSD semantics. The default is System V, which is recommended for most DMAs. This option is only applicable for locally attached tape drives exported via NDMP. Consult your DMA documentation for which mode your DMA expects. Changing this option only changes which devices are exported when the DMA scans for devices, so you will need to reconfigure the tape devices in your DMA after changing this setting.
DMA username and password	Used to authenticate the DMA. The system uses MD5 for user authentication

The NDMP service events log is available in `system-ndmpd:default`.

NFS Configuration

Network File System (NFS) is an industry-standard protocol to share files over a network. Oracle ZFS Storage Appliance supports NFS versions 2, 3, 4.0 and 4.1. For more information on how the filesystem namespace is constructed, see Working with Filesystem Namespace. For information about NFS with local users, see Configuring Users.

Caution:

To prevent loss of NFS service, as well as data loss, do not mount NFS filesystems using private interfaces.

To configure NFS, see the following sections:

- NFS Service Properties
- Configuring Kerberos Realms for NFS
- NFS Logs and Analytics
- NFS Properties
- NFS Naming Service Dependencies
- Sharing a Filesystem Over NFS

NFS Service Properties

The following NFS Service properties are available in **Configuration: Services**. Note that NFSv4 is also known as NFSv4.0.

- **Minimum supported version** - Use this drop-down list to control which versions of NFS the appliance supports.
- **Maximum supported version** - Use this drop-down list to control which versions of NFS the appliance supports.

  Note:

  Setting the NFS minimum and maximum versions to the same value causes the appliance to only communicate with clients using that version. This may be useful if you find an issue with one NFS version or the other (such as the performance characteristics of an NFS version with your workload), and you want to force clients to only use the version that works best.

- **Maximum # of server threads** - Define the maximum number of concurrent NFS requests (from 20 to 3000). This should at least cover the number of concurrent NFS clients that you anticipate. The default value is 1500.
- **Grace period** - Define the number of seconds that all clients have to recover locking state after an appliance reboot (from 15 to 600 seconds) from an unplanned outage. This property affects only NFSv4.0 and NFSv4.1 clients (NFSv3 is stateless so there is no state to reclaim). During this period, the NFS service only processes reclaims of the old locking state. No other requests for service are processed until the grace period is over. The default grace period is 90 seconds. Reducing the grace period lets NFS clients resume operation more quickly after a server reboot, but increases the probability that a client cannot recover all of its locking state. Oracle ZFS Storage Appliance provides grace-less recovery of the locking state for NFSv4.0 and NFSv4.1 clients during planned outages. Planned outages occur during events such as updates and appliance reboot using the CLI `maintenance system reboot` command or the BUI power icon ⏻. For planned outages, the NFS service processes all requests for service without incurring the grace period delay.

- **Custom NFSv4 identity domain** - Use this property to define the domain for mapping NFSv4.0 and NFSv4.1 users and group identities. If you do not set this property, the appliances uses DNS to obtain the identity domain, first by checking for a `_nfsv4idmapdomain` DNS resource record, and then by falling back to the DNS domain itself.
- **Use NFSv4 numeric id strings** - Use this property to allow NFSv4.0 and NFSv4.1 clients to use numeric strings for user and group IDs. If you do not set this property, user and group IDs are exchanged in the form of `user@domain`, the default. This property applies only when the authentication type is `AUTH_SYS`. The CLI property is `use_numeric_ids`.
- **Enable NFSv4 delegation** - Select this property to allow clients to cache files locally and make modifications without contacting the server. This option is enabled by default and typically results in better performance; but in rare circumstances it can cause problems. You should only disable this setting after careful performance measurements of your particular workload and after validating that the setting has a measurable performance benefit. This option only affects NFSv4.0 and NFSv4.1 mounts.
- **Mount visibility** - This property lets you limit the availability of information about share access lists and remote mounts from NFS clients. Full allows full access. Restricted restricts access such that a client can see only the shares which it is allowed to access. A client cannot see access lists for shares defined at the server or remote mounts from the server done by other clients. The property is set to `Full` by default.
- **Oracle Intelligent Storage Protocol** - The NFSv4.0 and NFSv4.1 services include support for the Oracle Intelligent Storage Protocol, which lets Oracle Database NFSv4.0 and NFSv4.1 clients pass optimization information to the Oracle ZFS Storage Appliance NFSv4.0 and NFSv4.1 server. For more information, see Configuring Oracle ZFS Storage Appliance for Oracle Database Clients.
- **Explicit netgroups** - If this property is false (default), the system applies heuristics to distinguish netgroups from hostnames in share access lists. Depending on the names in the access list and the responsiveness of DNS, these heuristics can result in a slow or unresponsive NFS service. If this property is true, netgroups are tagged (see section "NFS Protocol Share Mode Exceptions" in NFS Protocol) to distinguish them from hostnames, so the heuristics are no longer needed; specifically, no DNS lookups are performed to process netgroups. When setting this property to true, all netgroup names in a share access list must be prefixed with the `%` character in the CLI, or they must use the Netgroup exception type for the BUI. The preferred method to change this property, from either `true` to `false`, or `false` to `true`, is to use the `Netgroup editing workflow`. This applies the setting to all netgroup names in a share access list. For information about workflows, see Maintenance Workflows.

Related Topics

- NFS Properties
- Setting Service Properties - BUI, CLI

Configuring Kerberos Realms for NFS

Configuring a Kerberos realm creates certain service principals and adds the necessary keys to the system's local keytab. The NTP service must be configured before configuring Kerberized NFS. The following service principals are created and updated to support Kerberized NFS:

```
host/node1.example.com@EXAMPLE.COM
nfs/node1.example.com@EXAMPLE.COM
```

If you clustered your appliances, principals and keys are generated for each cluster node:

```
host/node1.example.com@EXAMPLE.COM
nfs/node1.example.com@EXAMPLE.COM
host/node2.example.com@EXAMPLE.COM
nfs/node2.example.com@EXAMPLE.COM
```

If these principals have already been created, configuring the realm resets the password for each of those principals.

For information on setting up KDCs and Kerberized clients, see Oracle Solaris documentation, which can be found at https://docs.oracle.com/en/operating-systems/solaris.html. For information about the appliance Kerberos service, see Kerberos Configuration. After configuring Kerberos, change the Security mode on the **Shares: Filesystem: Protocols** screen to a mode using Kerberos.

> **Note:**
>
> Kerberized NFS clients must access Oracle ZFS Storage Appliance using an IP address that resolves to an FQDN for those principals. For example, if an appliance is configured with multiple IP addresses, only the IP address that resolves to the appliance's FQDN can be used by its Kerberized NFS clients.

NFS Logs and Analytics

The following event logs are available for the NFS service.

Table 3-28 Logs Available for NFS

Log	Description
`network-nfs-server:default`	Master NFS server log
`appliance-kit-nfsconf:default`	Log of appliance NFS configuration events
`network-nfs-cbd:default`	Log for the NFSv4.0 and NFSv4.1 `callback` daemon
`network-nfs-mapid:default`	Log for the NFSv4.0 and NFSv4.1 `mapid` daemon, which maps NFSv4.0 and NFSv4.1 user and group credentials
`network-nfs-status:default`	Log for the NFS `statd` daemon, which assists crash and recovery functions for NFS locks
`network-nfs-nlockmgr:default`	Log for the NFS `lockd` daemon, which supports record locking operations for files

You can monitor NFS activity in the Analytics section. This includes:

- NFS operations per second
- ... by type of operation (read/write/...)
- ... by share name
- ... by client hostname
- ... by accessed filename
- ... by access latency

NFS Properties

The following table describes the mapping between CLI properties and the BUI properties.

> **Caution:**
>
> To prevent loss of NFS service, as well as data loss, do not mount NFS filesystems using private interfaces.

Table 3-29 NFS Properties

CLI Property	BUI Property
`version_min`	Minimum supported version
`version_max`	Maximum supported version
`nfsd_servers`	Maximum # of server threads
`grace_period`	Grace period
`mapid_domain`	Custom NFSv4 and NFSv4.1 identity domain
`use_numeric_ids`	Use NFSv4 and NFSv4.1 numeric string IDs
`enable_delegation`	Enable NFSv4 and NFSv4.1 delegation
`mount_visibility`	Client share information restriction level
`explicit_netgroups`	Use new syntax for netgroups in share access lists

NFS Naming Service Dependencies

Naming services, such as DNS, NIS, and LDAP, are used by Oracle ZFS Storage Appliance to resolve host names and corresponding IP addresses, user identities, and Analytics statistics. This topic describes NFS naming service dependencies and resulting problems if naming services are not configured for the appliance.

NFS depends on information from each of the naming services in the following table.

Table 3-30 NFS Naming Service Dependencies

Service	Description
DNS	IP address and corresponding host name of NFS clients and servers
NIS/LDAP	User identity number and corresponding user name
NIS/LDAP	Group identity number and corresponding group name
NIS/LDAP	Clients belonging to netgroups

If the appliance is unable to access any DNS network servers or DNS mappings are unpopulated, the following problems can occur:

- Filesystem mount failure
- Client is denied access to NFS shares after the filesystem has been successfully mounted
- Client receives "weak authentication" errors

- NFS server unresponsive
- User and group lookup failures, using either NFSv3 or NFSv4, as listed in the following table.

NFS Version	Problem	Setting
NFSv3 or NFSv4	Cannot access particular files or directories when the user is a member of 16 or more groups.	
NFSv4	Cannot retrieve ownership information (users and groups might be shown as `nobody`).	Set the option `Use NFSv4 numeric id strings`.
NFSv4	Cannot change ownership of files.	On the client, set the equivalent of option `Use NFSv4 numeric id strings`.
NFSv4	Cannot retrieve ACL information.	Set the option `Use NFSv4 numeric id strings`.
NFSv4	Cannot change ACLs, including inability to change entries unrelated to the affected entries.	Set the option `Use NFSv4 numeric id strings`. On the client, set the equivalent of option `Use NFSv4 numeric id strings`.

Related Topics

- DNS Configuration
- DNS-Less Operation

Sharing a Filesystem Over NFS

Use the following procedure to share a filesystem over NFS.

1. From the **Configuration** menu, select **Services**.
2. Check that the NFS service is enabled and online. If not, enable the service.

   > **Caution:**
   >
   > To prevent loss of NFS service, as well as data loss, do not mount NFS filesystems using private interfaces.

3. Click the **Shares** tab, and edit an existing share or create a new share.
4. Click the **Protocols** tab of the share you are editing, and check that NFS sharing is enabled.

   You can also configure the NFS share mode (`read/read+write`) in this screen.

NIS Configuration

Network Information Service (NIS) is a name service for centralized management. Oracle ZFS Storage Appliance can act as an NIS client for users and groups, with the following characteristics:

- NIS users can log in to the FTP and HTTP services.

- NIS users can be granted privileges for appliance administration. The appliance supplements NIS information with its own privilege settings.

To configure NIS, see the following sections:

- NIS Properties and Logs
- Adding an NIS User to the Appliance

NIS Properties and Logs

Table 3-31 NIS Properties

BUI Property	CLI Property	Description
Domain	`domain`	The NIS domain to use
Server(s): Search using broadcast	`broadcast`	Send an NIS broadcast to locate NIS servers for the specified domain
Server(s): Use listed servers	`ypservers`	Specify a list of NIS server hostnames or IP addresses

Oracle ZFS Storage Appliance connects to the first NIS server listed, or to a server found using broadcast, and switches to the next server if the current server stops responding.

Table 3-32 NIS Logs

Log	Description
`network-nis-client:default`	NIS client service log
`appliance-kit-nsswitch:default`	Log of the appliance name service, through which NIS queries are made
`system-identity:domain`	Log of the appliance domain name configurator

Related Topics

Adding an NIS User to the Appliance

Adding an NIS User to the Appliance

After you have completed NIS configuration, you can configure existing NIS users to be able to log in to Oracle ZFS Storage Appliance using their NIS username and password.

> **Note:**
>
> The existing NIS user ID cannot be less than 100, cannot be greater than 2147483646, and cannot be equal to 60001, 60002, or 65534. UIDs from 0-99 inclusive are reserved by the operating system vendor for use in future applications. Their use by end system users or vendors of layered products is not supported and can cause security issues with other applications.

In the BUI, follow the procedure Adding an Administrator or User (BUI), using the following property values:

1. Select **Directory** from the **Type** drop-down menu.
2. For **Username**, specify the user's existing NIS username.

In the CLI, follow the procedure Adding an Administrator or User (CLI), using the `directory` user type and the user's existing NIS username.

```
hostname:configuration users> directory NISusername
```

User ID (`uid`) and Password (`initial_password`) are set automatically: NIS values are used.

NTP Configuration

The Network Time Protocol (NTP) service can be used to keep the Oracle ZFS Storage Appliance clock accurate. This is important for recording accurate timestamps in the filesystem, and for protocol authentication. The appliance records times using the UTC timezone. The times that are displayed in the BUI use the timezone offset of your browser.

To the right of the BUI screen are times from both the appliance (**Server Time**) and your browser (**Client Time**). If the NTP service is not online, the **SYNC** button can be clicked to set the appliance time to match your client browser time.

If you are sharing filesystems using SMB, the client clocks must be synchronized to within five minutes of the appliance clock to avoid user authentication errors. One way to ensure clock synchronization is to configure the appliance and the SMB clients to use the same NTP server.

The NTP service events log is available in `network-ntp:default`.

To configure NTP, see the following sections:

- Setting Clock Synchronization (BUI)
- Configuring NTP (CLI)
- NTP Properties

Setting Clock Synchronization (BUI)

This procedure sets the Oracle ZFS Storage Appliance time to match the time of your web browser.

1. From the **Configuration** menu, select **Services**, then **NTP**.
2. Disable the NTP service.
3. Click **SYNC**.

Configuring NTP (CLI)

Use this procedure to configure NTP.

1. Go to `configuration services ntp` and enter `authkey` to edit authorizations.

   ```
   hostname:configuration services ntp> authkey
   hostname:configuration services ntp authkey>
   ```

2. From this context, new keys can be added with the `create` command.

   ```
   hostname:configuration services ntp authkey> create
   hostname:configuration services ntp authkey-000 (uncommitted)> get
   ```

```
                           keyno = (unset)
                            type = (unset)
                             key = (unset)
   hostname:configuration services ntp authkey-000 (uncommitted)> set keyno=1
                           keyno = 1 (uncommitted)
   hostname:configuration services ntp authkey-000 (uncommitted)> set type=A
                            type = SHA1 (uncommitted)
   hostname:configuration services ntp authkey-000 (uncommitted)> set key=coconuts
                             key = (set) (uncommitted)
   hostname:configuration services ntp authkey-000 (uncommitted)> commit
   hostname:configuration services ntp authkey>
   ```

3. To associate authentication keys with servers via the CLI, the `serverkeys` property should be set to a list of values in which each value is a key to be associated with the corresponding server in the servers property.

   If a server does not use authentication, the corresponding server key should be set to `0`. For example, use the following commands to use the key created earlier to authenticate the servers `server1` and `server2`:

   ```
   hostname:configuration services ntp> set servers=server1,server2
                         servers = server1,server2 (uncommitted)
   hostname:configuration services ntp> set serverkeys=1,1
                      serverkeys = 1,1 (uncommitted)
   hostname:configuration services ntp> commit
   hostname:configuration services ntp>
   ```

4. To associate authentication keys with servers, set the `serverkeys` property to a list of values in which each value is a key to be associated with the corresponding server in the `servers` property.

   If a server does not use authentication, the corresponding server key should be set to `0`. For example, use the following commands to use the key created earlier to authenticate the servers `server1` and `server2`.

   ```
   hostname:configuration services ntp> set servers=server1,server2
                         servers = server1,server2 (uncommitted)
   hostname:configuration services ntp> set serverkeys=1,1
                      serverkeys = 1,1 (uncommitted)
   hostname:configuration services ntp> commit
   hostname:configuration services ntp>
   ```

5. Use the following commands to authenticate the server `server1` with key `1`, `server2` with key `2`, and `server3` with key `3`.

   ```
   hostname:configuration services ntp> set servers=server1,server2,server3
                         servers = server1,server2,server3 (uncommitted)
   hostname:configuration services ntp> set serverkeys=1,2,3
                      serverkeys = 1,2,3 (uncommitted)
   hostname:configuration services ntp> commit
   hostname:configuration services ntp>
   ```

6. Use the following commands to authenticate the servers `server1` and `server2` with key `1`, and to additionally have an unauthenticated NTP server (`server3`).

   ```
   hostname:configuration services ntp> set servers=server1,server2,server3
                         servers = server1,server2,server3 (uncommitted)
   hostname:configuration services ntp> set serverkeys=1,1,0
                      serverkeys = 1,1,0 (uncommitted)
   hostname:configuration services ntp> commit
   hostname:configuration services ntp>
   ```

NTP Properties

The following NTP properties are available from **Configuration: Services: NTP**.

Table 3-33 NTP Properties

Property	Description	Examples
Discover NTP server via multicast address	Enter a multicast address here for an NTP server to be located automatically	`224.0.1.1`
Manually specify NTP server(s)	Enter one or more NTP servers (and their corresponding authentication keys, if any) for the appliance to contact directly	`0.pool.ntp.org`
NTP Authentication Keys	Enter one or more NTP authentication keys for the appliance to use when authenticating the validity of NTP servers. See the next table.	`Auth key: 10, Type: SHA225, Private Key: SUN7000`

Validation - If an invalid configuration is entered, a warning message is displayed and the configuration is not committed. This occurs when:

- A multicast address is used but no NTP response is found.
- An NTP server address is used, but that server does not respond properly to NTP.

Authentication - To prevent against NTP spoofing attacks from rogue servers, NTP has a private key encryption scheme whereby NTP servers are associated with a private key that is used by the client to verify their identity. These keys are not used to encrypt traffic, and they are not used to authenticate the client; they are only used by the NTP client (that is, Oracle ZFS Storage Appliance) to authenticate the NTP server. To associate a private key with an NTP server, the private key must first be specified. Each private key has a unique integer associated with it, along with a type and key. The type must be one of the following: RSA-SHA1, SHA1, SHA224, SHA256, SHA3-224, SHA3-256, SHA3-384, SHA3-512, SHA384, SHA512, SHA512-224, SHA512-256, SHAKE128, and SHAKE256. Each type is a hash algorithm and takes a key value that is an ASCII string. To authenticate the response, the key number, key type, and key value on the client must match the key number, key type, and key value on the server.

After the keys have been specified, an NTP server can be associated with a particular private key. For a given key, all of the key number, key type, and private key values must match between client and server for an NTP server to be authenticated.

Phone Home Configuration

The **Phone Home** service screen is used to manage the Oracle ZFS Storage Appliance registration, as well as the Phone Home remote support service.

Registration connects your appliance with Oracle Auto Service Request (https://www.oracle.com/support/premier/auto-service-request.html). Oracle ASR automatically opens Service Requests (SR) for specific problems reported by your appliance. Registration also connects your appliance with My Oracle Support to detect update notifications.

The Phone Home service communicates with Oracle Support to provide:

- **Fault reporting** - The system reports active problems to Oracle for automated service response. Depending on the nature of the fault, a support case may be opened. Details of these events can be viewed in the Active Problem Display. For more information, see Working with Problems in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

- **Heartbeats** - Daily heartbeat messages are sent to Oracle to indicate that the system is up and running. Oracle Support may notify the technical contact for an account when one of the activated systems fails to send a heartbeat for too long.
- **System configuration** - Periodic messages are sent to Oracle describing current software and hardware versions and configuration as well as storage configuration. No user data or metadata is transmitted in these messages.
- **Support bundles** - The Phone Home service must be enabled before support bundles can be uploaded to Oracle Support. See Working with Support Bundles in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x* for more information.
- **Update notifications** - Creates an alert when new software updates are available on My Oracle Support (MOS). See Working with Software Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x* for more information.

You must register to use the Phone Home service. You need a valid Oracle Single Sign-On account user name and password to use the fault reporting and heartbeat features of the Phone Home service. Go to My Oracle Support (https://support.oracle.com) and click **Register** to create your account.

To configure Phone Home, see the following sections:

- Registering the Appliance - BUI, CLI
- Changing Account Information (BUI)
- Phone Home Properties

Registering the Appliance (BUI)

Use the following procedure to register your Oracle ZFS Storage Appliance with Oracle Support.

1. From the **Configuration** menu, select **Services**, then **Phone Home**.
2. Enter your Oracle Single Sign-On Account user name and password.

   Click **Privacy Statement** for information about privacy policy. It can be viewed at any time in both the BUI and CLI.
3. Click **APPLY** to commit your changes.
4. Use My Oracle Support (https://support.oracle.com/) to complete activation for Oracle Auto Service Request (https://www.oracle.com/support/premier/auto-service-request.html).

   Refer to "How To Manage and Approve Pending ASR Assets In My Oracle Support" (Doc ID 1329200.1).

Registering the Appliance (CLI)

Use the following procedure to register your Oracle ZFS Storage Appliance with Oracle Support.

1. Go to `configuration services scrk`.
2. Set `soa_id` and `soa_password` to the user name and password for your Oracle Single Sign-On Account, respectively.
3. Commit your changes.

4. Use My Oracle Support (https://support.oracle.com/) to complete activation for Oracle Auto Service Request (https://www.oracle.com/support/premier/auto-service-request.html).

   Refer to "How To Manage and Approve Pending ASR Assets In My Oracle Support" (Doc ID 1329200.1).

Example 3-3 CLI Registration

```
hostname:>
          configuration services scrk

hostname:configuration services scrk>
          set soa_id=myuser

                         soa_id = myuser(uncommitted)
hostname:configuration services scrk>
          set soa_password=mypass

                   soa_password = (set) (uncommitted)
hostname:configuration services scrk>
          commit
```

Changing Account Information (BUI)

Use the following procedure to change account information with Oracle Support for your Oracle ZFS Storage Appliance.

1. From the **Configuration** menu, select **Services**, then **Phone Home**.
2. Click **Change account** to change the Oracle Single Sign-On Account used by the appliance.
3. Commit your changes.
4. Use My Oracle Support to complete Auto Service Request (ASR) activation.

   Refer to "How To Manage and Approve Pending ASR Assets In My Oracle Support" (Doc ID 1329200.1).

Phone Home Properties

If Oracle ZFS Storage Appliance is not directly connected to the Internet, you might need to configure an HTTP proxy through which the Phone Home service can communicate with Oracle. These proxy settings will also be used to upload support bundles. For more details on support bundles, see Working with Support Bundles in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

Table 3-34 Phone Home Web Proxy Settings

Property	Description
Use web proxy	Connect via a web proxy
Host : port	Web proxy hostname or IP address, and port
Username	Web proxy username
Password	Web proxy password

Table 3-35 Phone Home Status

Property	Description
Last heartbeat sent at	Time last heartbeat was sent to Oracle Support

If the Phone Home service is enabled before a valid Oracle Single Sign-On account has been entered, it will appear in the maintenance state. You must enter a valid Oracle Single Sign-On account to use the Phone Home service.

There is a log of Phone Home events in **Maintenance: Logs: Phone Home**.

RADIUS Configuration

> **Caution:**
>
> When the appliance RADIUS service is enabled, *all* directory users log in using RADIUS. To create a directory user, see Configuring Users.

Oracle ZFS Storage Appliance supports the RADIUS (Remote Authentication Dial-In User Service) directory service for centralized authentication of directory users.

The RADIUS service is a client-server protocol used in conjunction with a RADIUS server to authenticate directory users for logging in to remote systems, such as Oracle ZFS Storage Appliance.

Much like the appliance's LDAP service, the RADIUS service communicates with another server that contains the user database.

RADIUS provides authentication by supporting classic password-based authentication, as well as supporting multi-factor authentication, which requires additional authentication using such schemes as challenge-response authentication and one-time password authentication. Multi-factor authentication adds a layer of security to help prevent unauthorized access.

The RADIUS server, not Oracle ZFS Storage Appliance, controls the authentication process and controls the prompts for required information.

The Oracle ZFS Storage Appliance RESTful API supports RADIUS authentication when only a single response is required, such as a password. Authentication sequences requiring multiple prompts and responses, such as a password, a challenge, and a response to the challenge, are not supported.

As detailed in the following topics, this section describes how to configure Oracle ZFS Storage Appliance for use with RADIUS servers, and how to monitor the RADIUS servers from the appliance:

- RADIUS Properties and Logs
- Configuring RADIUS Servers - BUI, CLI
- Configuring RADIUS Server Certificates
- Monitoring RADIUS Server Status - BUI, CLI

RADIUS Properties and Logs

This section contains the RADIUS general settings properties, servers properties, and the RADIUS service events log location.

The following table lists the BUI and CLI RADIUS general settings properties.

Table 3-36 RADIUS General Settings Properties

BUI Property	CLI Property	Description
Transport Protocol • Require TLS • Unencrypted UDP (not recommended)	`protocol` • `tls` • `udp`	Transport protocol between Oracle ZFS Storage Appliance and the RADIUS server. For a description of these choices, see Configuring RADIUS Servers - BUI, CLI.
Server Shared Secret	`secret`	For UDP connections, this is the shared secret. When this property is not set, it is blank. When this property is set, it is shown as `(set)`.
RADIUS Servers	`servers`	List of RADIUS servers, with each server defined as `host[:port]`. `host` can be a host name, an IPv4 address, or an IPv6 address. The `port` parameter is optional. If an IPv6 port is specified in the CLI, surround the IPv6 address with brackets; for example, `[2001:db8:3c4d:0015:0000:0000:1a2f:1a2b]:1812`.

The following table lists the BUI and CLI RADIUS server properties.

Table 3-37 RADIUS Server Properties

BUI Property	CLI Property	Description
RTT	`rtt`	Round-trip time for Oracle ZFS Storage Appliance to receive a response from the RADIUS server. In the list of servers, if the server cannot be contacted, the RTT is replaced by an error message.
	`err_msg`	If the server cannot be contacted, the error message provides the reason.

The RADIUS service events log can be viewed in the BUI as described in Monitoring RADIUS Server Status (BUI).

Related Topics

- Configuring RADIUS Server Certificates
- Configuring RADIUS Servers - BUI, CLI
- Monitoring RADIUS Server Status - BUI, CLI

Configuring RADIUS Servers (BUI)

Use the following procedure to configure Oracle ZFS Storage Appliance for use with RADIUS servers, and then enable the RADIUS service.

> **Caution:**
>
> Enabling the TLS protocol is highly recommended so that critical information is sent securely.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **RADIUS**.
3. Select a transport protocol.
   - **Require TLS** – Use the TLS protocol to securely connect to the RADIUS server.
   - **Unencrypted UDP (not recommended)** – Use the UDP protocol to connect to the RADIUS server. The UDP protocol is not considered as secure as the TLS protocol.
   - **Server Shared Secret** – When adding a server that uses a UDP connection, enter the shared secret.
4. Configure the list of RADIUS servers.
   - **Add Server:**
     a. Click the add icon ⊕ to the left of **RADIUS Servers**.
     b. In the **New RADIUS Server** dialog box, enter the server name or IP address and optionally the port number (after the colon). The server name can be a host name, such as `*hostname*.example.com`, an IPv4 address, or an IPv6 address.
     c. For a UDP connection, click **APPLY**.
     d. For a TLS connection, click **APPLY**. One of two dialog boxes is displayed: Either a dialog box confirms that the certificate is trusted and you click **OK**, or the **Accept Server Certificate** dialog box opens. Click either **ACCEPT** or **REJECT** for the certificate details. If you accept the certificate, it is added to the list of trusted certificates. For more information, see Configuring RADIUS Server Certificates.
     e. Click **APPLY** at the top of the RADIUS properties page.
   - **Change Server Name, Port Number, IP Address:**
     a. Click on the row of the server that you want to modify, and click the edit icon ✎ to the right of the row.
     b. In the **Edit RADIUS Server** dialog box, change the server name, port number or IP address.
     c. (Optional) Click **Test Connection**.
     d. When you are finished, click **APPLY** in the **Edit RADIUS Server** dialog box.
     e. Click **APPLY** at the top of the RADIUS properties page.
   - **Remove Server:**
     a. Click on the row of the server that you want to remove, and click the trash icon 🗑 to the right of the row.

   b. When finished, click **APPLY** at the top of the RADIUS properties page.

5. Enable the RADIUS service.

   a. From the **Configuration** menu, select **Services**.

   b. Click the power icon ⏻ for the RADIUS service to bring the service online .

Related Topics

- RADIUS Properties and Logs
- Configuring RADIUS Server Certificates
- Monitoring RADIUS Server Status (BUI)

Configuring RADIUS Servers (CLI)

Use the following procedure to configure Oracle ZFS Storage Appliance for use with RADIUS servers, and then enable the RADIUS service.

> **Caution:**
>
> Enabling the TLS protocol is highly recommended so that critical information is sent securely.

1. Go to `configuration services radius`.

2. Enter the `show` command to show the RADIUS service properties:

```
hostname:configuration services radius> show
Properties:
    <status> = offline
     servers =
    protocol =
      secret =
```

3. Set the transport protocol and commit the change:

```
hostname:configuration services radius> set protocol=tls
    protocol = tls (uncommitted)
hostname:configuration services radius> commit
```

   When setting the UDP protocol, set the shared secret and commit the change.

4. Configure the list of RADIUS servers.

   - **Add Server:**

     a. Create the server, set the server’s name or IP address and optional port number, and commit the change. If an IPv6 port is specified, surround the IPv6 address with brackets.

```
hostname:configuration services radius> create
hostname:configuration services radius (uncommitted)> set server =
hostname.example.com
    server = hostname.example.com (uncommitted)
hostname:configuration services radius> commit
```

     b. For a TLS connection, if the certificate is not already trusted, you are prompted to accept the certificate. Enter **`y`** or **`n`**, as appropriate.

For a UDP connection, if there is a problem contacting the service, you are prompted to confirm the settings. Enter **y** or **n**, as appropriate.

- **Change Server Name, Port Number, IP Address:**

  a. List the servers to display the servers' ordinal names:

  ```
  hostname:configuration services radius> list
  SERVER     STATUS      RADIUS SERVER
  server-000 online      hostname.example.com
  server-001 unavailable host.sample.com
  ```

  b. Select a server by its ordinal name:

  ```
  hostname:configuration services radius> select server-000
  hostname:configuration services server-000>
  ```

  c. To change the server's name and port number or the IP address, set the `server` property and commit the change:

  ```
  hostname:configuration services server-000> set server=name.sample.com:484
      server = name.sample.com:484 (uncommitted)
  hostname:configuration services server-000> commit
  ```

  The same errors could occur as when adding a new server.

- **Remove Server:**

  a. List the servers to display the servers' ordinal names:

  ```
  hostname:configuration services radius> list
  SERVER     STATUS      RADIUS SERVER
  server-000 online      hostname.example.com
  server-001 unavailable host.sample.com
  ```

  b. To remove the server, enter the `destroy` command, followed by the server's ordinal name:

  ```
  hostname:configuration services server-001> destroy server-001
  ```

  c. Confirm your action.

5. Optional: Test a server's TLS connection.

   a. List the servers to display the servers' ordinal names:

   ```
   hostname:configuration services radius> list
   SERVER     STATUS      RADIUS SERVER
   server-000 online      hostname.example.com
   server-001 unavailable host.sample.com
   ```

   b. Select a server by its ordinal name:

   ```
   hostname:configuration services radius> select server-001
   hostname:configuration services server-001>
   ```

   c. Enter command `test`:

   ```
   hostname:configuration services server-001> test
   ```

6. Enable the RADIUS service.

   a. Go to `configuration services radius`.

   b. Enter command `enable` to bring the service online:

   ```
   hostname:configuration services radius> enable
   ```

Related Topics

- RADIUS Properties and Logs
- Configuring RADIUS Server Certificates
- Monitoring RADIUS Server Status (CLI)

Configuring RADIUS Server Certificates

If the RADIUS service uses a TLS connection, a valid RADIUS server certificate must be used with the service. A RADIUS server's certificate can be CA-signed or self-signed. This section describes how to initially configure certificates and how to manage a new certificate when the previous certificate expires. This section does not apply to the UDP protocol.

Initially Configuring RADIUS Server Certificates

For more information about trusted certificates, see the sections about trusted certificates in Configuring Certificates.

You can supply a list of trusted CA certificates. RADIUS server certificates issued by those trusted CAs and marked as trusted by RADIUS do not require special management.

If a RADIUS server's certificate is not issued by a trusted CA, whether the certificate is issued by a CA or is self-signed, you will be asked to review and approve the certificate. If you accept the certificate, that certificate is added to the list of trusted certificates.

Managing Expired and New RADIUS Server Certificates

If you individually accepted a certificate, either a CA-signed certificate or a self-signed certificate, then when the RADIUS server's certificate expires, you must approve the new certificate. Select the server, test the connection, and examine and approve the new certificate. See Approving a New RADIUS Server Certificate - BUI, CLI.

If you supply CA certificates, changes in the individual server certificates are handled automatically. When your server changes CA certificates, ensure that the new CA certificate is added to the appliance before your RADIUS server starts using it. If the server starts using the new CA certificate before you add it to the appliance, your RADIUS service will be interrupted.

Approving a New RADIUS Server Certificate (BUI)

Use the following procedure to accept a new certificate after the previous certificate has expired.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **RADIUS**.
3. On the **Properties** tab, scroll to the **RADIUS Servers** section of the page.
4. In the table, click the edit icon ✎ for the server that has the new certificate.
5. Click the **Test Connection** button in the **Edit RADIUS Server** dialog box to test the TLS connection.

   A new dialog box reports whether the new certificate is trusted.
6. Click **OK** in the trusted certificate dialog box.

   If the trusted certificate dialog box reported that the certificate is not trusted, the **Accept RADIUS Server Certificate** dialog box opens. This dialog box displays information about the certificate, and has **REJECT** and **ACCEPT** buttons.
7. Review the certificate information, and click **ACCEPT**.

The certificate is added to the list of trusted certificates.

Related Topics

Viewing Trusted Certificate Details (BUI)

Approving a New RADIUS Server Certificate (CLI)

Use the following procedure to accept a new certificate after the previous certificate expired.

1. Go to `configuration services radius`.
2. Enter the `list` command to show the list of RADIUS servers' ordinal names.
3. Select a server by its ordinal name.
4. Enter the `test` command to test the TLS connection.

   Information about the new certificate is displayed.
5. Examine and approve the new certificate.

   The certificate is added to the list of trusted certificates. If you enter the `test` command again, the message `Certificate is trusted` is displayed.

Related Topics

Viewing Trusted Certificate Details (CLI)

Monitoring RADIUS Server Status (BUI)

Use the following procedure to monitor the RADIUS server status.

1. From the **Configuration** menu, select **Services**.
2. Under **Directory Services**, select **RADIUS**.
3. View server status.

   On the **Properties** tab, scroll to the **RADIUS Servers** section of the page.

   For each server, the following status is displayed:

   - **Status icon -** Indicator icon that represents the status of the server. The indicator icon is either online or unavailable ● .
   - **RTT** - The round-trip time for the appliance to get a response from the RADIUS server. If the server cannot be contacted, an error message is shown in its place.
4. View the RADIUS service logs.
   a. Click the **Logs** tab at the top of the RADIUS page.
   b. If there are multiple pages, use the double arrow keys to navigate between the pages. To go to the first page, click the left-facing single arrow key. To go to the last page, click the right-facing single arrow key.

Related Topics

- RADIUS Properties and Logs
- Configuring RADIUS Servers (BUI)

Monitoring RADIUS Server Status (CLI)

Use the following procedure to monitor the RADIUS server status.

1. Go to `configuration services radius`.
2. Enter the `show` command to show RADIUS service properties, as well as the RADIUS servers listed by their ordinal names and the round-trip time (RTT) or error message.

```
hostname:configuration services radius> show
Properties:
    <status> = online
     servers = hostname.example.com,host.sample.com
    protocol = tls
      secret = (set)

Servers:
SERVER     RADIUS SERVER        RTT
server-000 hostname.example.com  0.680ms
server-001 host.sample.com       Connection refused
```

3. To view the service logs, use the BUI, as described in Monitoring RADIUS Server Status (BUI).

Related Topics

- RADIUS Properties and Logs
- Configuring RADIUS Servers (CLI)

RESTful API Configuration

The Oracle ZFS Storage Appliance RESTful API lets you manage the appliance using simple requests such as `GET`, `PUT`, `POST`, and `DELETE HTTP` against managed resource URL paths.

The appliance RESTful-based architecture is defined as a layered client-server model. Advantages of this model mean that services can be transparently redirected through standard hubs, routers, and other network systems without client configuration. This architecture supports caching of information and is useful when many clients request the same static resources.

For complete Oracle ZFS Storage Appliance RESTful API documentation, see the Oracle ZFS Storage Appliance RESTful API Guide, Release OS8.8.x.

Service Tags Configuration

Service tags are used to facilitate product inventory and support, by allowing Oracle ZFS Storage Appliance to be queried for data such as:

- System serial number
- System type
- Software version numbers

You can register the service tags with Oracle Support, allowing you to easily keep track of your Oracle equipment and also expedite service calls. The service tags are enabled by default.

Table 3-38 UDP/TCP Port Properties

Property	Description
Discovery Port	UDP port used for service tag discovery; default is 6481
Listener Port	TCP port used to query service tag data; default is 6481

SFTP Configuration

The SFTP (SSH File Transfer Protocol) service allows filesystem access from SFTP clients. Anonymous logins are not allowed, users must authenticate with whichever name service is configured in **Services**.

SFTP can be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see Kerberos Configuration.

For added security when configuring SFTP, you can specify the ciphers and MACs, as described in SFTP Properties, Ports, and Logs.

To configure SFTP, see the following sections:

- Adding SFTP Access to a Share (BUI)
- Configuring SFTP for Remote Access (CLI)
- SFTP Properties, Ports, and Logs

Adding SFTP Access to a Share (BUI)

Use the following procedure to add SFTP access to a share.

1. From the **Configuration** menu, select **Services**.
2. Check that the SFTP service is enabled and online. If not, enable the service.
3. From the **Shares** menu, select **Shares**, and select or add a share.
4. Go to the **Protocols** tab, and check that SFTP access is enabled.
5. Optional: Set the **Share mode access** to **Read only** or **Read/Write**.

Related Topics

- Configuring SFTP for Remote Access (CLI)
- SFTP Properties, Ports, and Logs

Configuring SFTP for Remote Access (CLI)

Use the following procedure to configure SFTP for remote access.

1. Create a local user or network user (LDAP or NIS) with an appropriate administrator role. (See Configuring Users.)
2. Generate an SSH authentication key by entering the command `ssh-keygen -t dsa` on the Oracle Solaris host/client.
3. Enter a file name in which to store the key.
4. Enter a passphrase if required, or leave this field blank to log on directly to the SFTP share.

The location is displayed for the key. The key looks similar to the following:

```
ssh-dss AAAAB3NzaC1kc3MAAACBAPMMs5h8UWk1NPf/VJDDEo0OAwT+s6iZxkCmmrgAmLfTX9izWk+
bsvNldOlXN/6EgkusLjo/+UaEt5+704vMHClRaq3AlVHLS5tVjeX3iCs+fDo0qwXZg3Brh8QBAaWk3
ywr2osuII1tHh4v/HwEAHZq5mVWXav0pO3bgmxl0/+VAAAAFQDIJxnm52DfyEdQQMTY+jRVvzGwMQA
AAIAhTP6Ey+2gGFiCKkvUofsco4d8pbqH8duE9P6Y88s0+opuj52GkAdRUt2fRrdM9Cf3h4lIOc8Bw9
bZIBzrCKBNWBUdZG56tsfLdilW6vS6gxKrmL2v7fSp9WYPsxZGhOLfU29zW4n2WVcVHbGyFEoVe+taq
aq+AYJaWoHnjZL1/LpQAAAIAOLc8+uc3hDOcK3pAkYdg8b2rYIGOAZU4py0rq24DGPeVHd5h5jbe4p
WDM70uYqGCOPYiOKeEoMNJpczRX5qjI+BfoUY4sH24WWwsKkT8XX9PUAa0WT+7axEqg2N6YelaTJ95J
vMaj6E7HkAIra2Sj2H/LSDktL42UL+j1Wx5A==username sunray
```

5. From the **Configuration** menu, select **Services**, then **SFTP**. Under **Keys**, click the plus (+) sign.
6. In the **New Key** window, select **DSA**.
7. Copy only the key portion (in the previous example, the key begins with `AAAA` and ends with `Wx5A==`), and paste it into the **Key** field.

   > **Note:**
   >
   > The key should not contain any white (negative) spaces.

8. Enter the user name, and add a comment as a reminder.
9. From the **Shares** menu, select **Shares**, and click the add item icon ⊕ to create a filesystem.
10. In the **Create Filesystem** window, enter the filesystem name (for example, `sftp`), change the permissions to **Read/ Write** for the share, and click **APPLY**.
11. Click the edit icon ✎ to set up the share properties. (See Filesystem Properties.)
12. To access the share, use the `sftp` command as shown in these examples:

    ```
    sftp -o "port=218" username 10.x.x.151:/export/sftp
    Connecting to 10.x.xx.151...
    Changing to: /export/sftp
    sftp>
    ```

    Example with `-v` option:

    ```
    sftp -v -o "IdentityFile=/home/username/.ssh/id_dsa" -o "port=218"
    root 10.x.xx.151:/export/sftp
    ```

Related Topics

- Adding SFTP Access to a Share (BUI)
- SFTP Properties, Ports, and Logs

SFTP Properties, Ports, and Logs

The following tables list the SFTP properties, security properties, ports, and logs.

SFTP Properties

Table 3-39 SFTP Properties

Property	Description
Port (for incoming connections)	The port SFTP listens on. The default is 218.
Permit root login	Allows SFTP logins for the root user. This property is off by default.
Logging level	The verbosity of SFTP log messages.
Idle Session Timeout	Idle timeout in seconds for client session. After the timeout value has been reached and if there is no activity, the user session is closed. By default, the value is set to `Infinite`.
Keys	RSA/DSA public keys for SFTP authentication. Text comments can be associated with the keys to help administrators track why they were added. As of the 2011.1 software release, key management for SFTP has changed to increase security. When creating an SFTP key, it is required to include the `user` property with a valid user assignment. SFTP keys are grouped by user and are authenticated via SFTP with the user's name. It is recommended to recreate any existing SFTP keys that do not include the user property, even though they will still authenticate.

Table 3-40 SFTP Security Properties

Property	Description
Ciphers	Ciphers for SFTP connections.
MACs	Message authentication codes (MACs) for SFTP connections.

SFTP Ports

The SFTP service uses a non-standard port number for connections to Oracle ZFS Storage Appliance. This is to avoid conflicts with administrative SSH connections to port 22. By default, the SFTP port is 218 and must be specified on the SFTP client prior to connecting. For example, an Oracle Solaris client using SFTP would connect with the following command:

```
manta# sftp -o "Port 218" root@hostname
```

SFTP Logs

The SFTP service events log is available at `network-sftp:default`.

Related Topics

- Adding SFTP Access to a Share (BUI)
- Configuring SFTP for Remote Access (CLI)

Shadow Migration Configuration

The Shadow Migration service allows for automatic migration of data from external or internal sources. This functionality is described in detail in Shadow Migration. The service itself only controls automatic background migration. Regardless of whether the service is enabled or not, data will be migrated synchronously for in-band requests.

The service should only be disabled for testing purposes, or if the load on the system due to shadow migration is too great. When disabled, no filesystems will ever finish migrating. The primary purpose of the service is to allow tuning of the number of threads dedicated to background migration.

Number of Threads Property - Number of threads to devote to background migration of data. These threads are global to the entire machine, and increasing the number can increase concurrency and the overall speed of migration at the expense of increased resource consumption (network, I/O, and CPU).

SMB Configuration

The SMB service provides access to filesystems using the SMB protocol. The supported SMB versions are: SMB 1, SMB 2.0, SMB 2.1, SMB 3.0, and SMB 3.1. To share filesystems over SMB, configure the filesystem as described in Filesystem Properties. The following tables show the supported and unsupported features for SMB 3.1, SMB 3.0, and SMB 2.1.

It is strongly advised to upgrade clients from SMB 1 to at least SMB 2.0 because SMB 1 has known security and performance issues that are resolved in later SMB versions.

Table 3-41 SMB 3.1 Supported and Unsupported Features

Supported Features	Unsupported Features
Pre-authentication integrity	
Encryption improvements: Added AES-128-GCM	

Table 3-42 SMB 3.0 Supported and Unsupported Features

Supported Features	Unsupported Features
Transparent failover (Continuously Available shares)	SMB over Remote Direct Memory Access (RDMA)
Multichannel	Volume Shadow Copy Service (VSS) for SMB filesystems
Encryption	Directory leasing

Table 3-43 SMB 2.1 Supported and Unsupported Features

Supported Features	Unsupported Features
Lease	Branch cache
Multi-protocol negotiate request	Resilient handles
Individual write-through operations	
Multi-credit operations	

Local accounts and user IDs are mapped to Windows user IDs. Note that the *guest* account is a special, read-only account and cannot be configured for read/write in Oracle ZFS Storage Appliance.

To configure SMB, see the following sections:

- SMB Service Properties
- Setting Properties to Export Shares over SMB
- NFS/SMB Interoperability

- SMB DFS Namespaces
- SMB Microsoft Stand-alone DFS Namespace Management Tools Support Matrix
- Adding DFS Namespaces to a Local SMB Group
- SMB Autohome
- Adding SMB Autohome Rules (CLI)
- Adding a User to an SMB Local Group
- SMB MMC Integration
- SMB Share Management
- SMB Users, Groups, and Connections
- Listing SMB Services
- Configuring SMB (BUI)
- Configuring SMB Active Directory (BUI)
- Configuring SMB Project and Share (BUI)
- Configuring SMB Data Service (BUI)

SMB Service Properties

Changing service properties is documented in Setting Service Properties (BUI) and Setting Service Properties (CLI).

- **Minimum supported version** - Choose the minimum version of SMB that Oracle ZFS Storage Appliance supports.
- **Maximum supported version** - Choose the maximum version of SMB that the appliance supports.
- **System comment** - Meaningful text string.
- **Idle Session timeout** - Timeout setting for session inactivity.
- **Preferred domain controller** - The preferred domain controller to use when joining an Active Directory domain. If this controller is not available, Active Directory will rely on DNS SRV records and the Active Directory site to locate an appropriate domain controller. For more information, see Active Directory Configuration.
- **Active Directory site** - The site to use when joining an Active Directory domain. A site is a logical collection of machines which are all connected with high bandwidth, low latency network links. When this property is configured and the preferred domain controller is not specified, joining an Active Directory domain will prefer domain controllers located in this site over external domain controllers.
- **Maximum # of server threads** - The maximum number of simultaneous server threads (workers). Default is 1024.
- **Enable Dynamic DNS** - Choose whether the appliance will use Dynamic DNS to update DNS records in the Active Directory domain. Default is off.
- **Enable oplocks** - Choose whether the appliance will grant opportunistic locks to SMB clients. This will improve performance for most clients. Default is on. The SMB server grants an oplock to a client process so that the client can cache data while the lock is in place. When the server revokes the oplock, the client flushes its cached data to the server.

- **Restrict anonymous access to share list** - If this option is enabled, clients must authenticate to the SMB service before receiving a list of shares. If disabled, anonymous clients may access the list of shares.
- **Primary WINS server** - Primary WINS address configured in the TCP/IP setup.
- **Secondary WINS server** - Secondary WINS address configured in the TCP/IP setup.
- **Excluded IP addresses from WINS** - IP addresses excluded from registration with WINS.
- **LAN Manager compatibility level** - Authentication modes supported (LM, NTLM, LMv2, NTLMv2). For more information on the supported authentication modes within each compatibility level, consult the Oracle Solaris Information Library for *smb*. NTLMv2 is the recommended minimum security level to avoid publicly known security vulnerabilities.
- **SMB signing enabled** - Enables interoperability with SMB clients using the SMB signing feature. If a packet has been signed, the signature will be verified. If a packet has not been signed it will be accepted without signature verification (if SMB signing is not required, see below).
- **SMB signing required** - When SMB signing is required, all SMB packets must be signed or they will be rejected, and clients that do not support signing will be unable to connect to the server.
- **Ignore zero VC** - When an SMB client establishes a new connection, it may request that the appliance clean up all previous connections and file locks from this client by specifying a Virtual Circuit (VC) number of zero. This protocol artifact however, does not respect network address translation (NAT) for clients or multiple DNS entries assigned to the same host. In combination, zero VC requests between masked or redundant network locations may result in unrelated active connections being reset. By default, zero VC requests are honored to prevent stale file locking, however if SMB sessions are being disconnected in error, ignoring zero VC requests may resolve the issue.
- **Share visibility** - Use this property to set the access-based enumeration (ABE) policy for displaying available shares to clients. Valid values are `Full` and `Restricted`. While `Full` allows full access, `Restricted` limits access to only shares that the client is allowed to see. Access to shares is determined by the SMB exceptions and the share's ACL. This property is set to `Full` by default.
- **NetBIOS enable** - Enables or disables all NetBIOS services. A value of true (default) enables NetBIOS name (UDP port 137), datagram (UDP port 138), and session (TCP port 139) services, and enables locating the domain controller via NetBIOS-based discovery, while a value of false disables all of them.
- **Encrypt data access** - Enables the SMB server to require that clients encrypt data on all new sessions. This enforcement can be bypassed if the server allows unencrypted access. This configures SMB encryption at the global level, and the default value is `false`. See also "Reject unencrypted access."
- **Reject unencrypted access** - Rejects unencrypted access when either global-level encryption or share-level encryption is enabled. The default value is `true`. When set to `false`, unencrypted access is allowed. Do not set this property to `false` unless security implications are understood. Allowing unencrypted access might be acceptable when a deployment scenario requires support to down-level clients that do not support encryption.
- **Enable multi-channel** - Enables or disables SMB3 multi-channel support. When set to `true`, the default, the SMB server accepts multi-channel paths between the SMB server and client. Disabling multi-channel support could be beneficial for some firewall configurations. Use the `multichannel_exclude` property to specify physical interfaces that are not to be used for SMB multi-channel. Private and deprecated interfaces are automatically excluded.

- **Explicit netgroups** - If this property is `false` (default), the system applies heuristics to distinguish netgroups from hostnames in share access lists. Depending on the names in the access list and the responsiveness of DNS, these heuristics can result in a slow or unresponsive SMB service. If this property is `true`, netgroups are tagged (see section SMB Protocol Share Mode Exceptions) to distinguish them from hostnames, so the heuristics are no longer needed; specifically, no DNS lookups are performed to process netgroups. When setting this property to `true`, all netgroup names in a share access list must be prefixed with the `%` character in the CLI or they must use the **Netgroup exception type** for the BUI. The preferred method to change this property, from either `true` to `false`, or `false` to `true`, is to use the "Netgroup editing workflow." This applies the setting to all netgroup names in a share access list. For information about workflows, see Maintenance Workflows.
- **Maximum machine account password age** - Specifies the number of days, from 1 to 999, until the next Active Directory computer account password change. This property is applicable only when Oracle ZFS Storage Appliance is joined to an Active Directory domain. To disable this property in the BUI, select **Disable periodic password change**. To disable in the CLI, set the property to 0. This property is disabled by default.

  It is recommended to set the value to 30 days. Values lower than this can increase replication efforts and affect domain controllers. Significantly increasing the value or disabling the property gives an attacker more time to undertake a brute-force password-guessing attack against one of the machine accounts.

Setting Properties to Export Shares over SMB

Several share properties must be set in certain ways when exporting a share over SMB.

Table 3-44 SMB Share Properties

Property	Description
Case Sensitivity	SMB clients expect case-insensitive behavior, so this property must be `mixed` or `insensitive`. See Static Properties.
Reject non UTF-8	If non-UTF-8 filenames are allowed in a filesystem, SMB clients may function incorrectly. See Static Properties.
Non-Blocking Mandatory Locking	This property must be enabled to allow byte range locking to function correctly. See Static Properties.
Resource name	The name by which clients refer to the share. For information about how this name is inherited from a project, see Share and Project Protocols.
Share-level ACL	An ACL which adds another layer of access control beyond the ACLs stored in the filesystem. For more information on this property, see Access Control Lists for Filesystems.

The case sensitivity and reject non UTF-8 properties can only be set when creating a share.

No two SMB shares on the same system may share the same resource name. Resource names inherited from projects have special behavior. For details, see Shares and Projects. Resource names must be less than 80 characters, and can contain any alphanumeric characters besides the following characters:

```
" / \ [ ] : | < > + ; , ? * =
```

When access-based enumeration is enabled, clients may see directory entries for files which they cannot open. Directory entries are filtered only when the client has no access to that file.

For example, if a client attempts to open a file for read/write access but the ACL grants only read access, that open request will fail but that file will still be included in the list of entries.

NFS/SMB Interoperability

Oracle ZFS Storage Appliance supports NFS and SMB clients accessing the same shares concurrently. To correctly configure the appliance for NFS/SMB interoperability, you must configure the following components:

- Configure the Active Directory service. See Active Directory Configuration.
- Establish an identity mapping strategy and configure the service. See Identity Mapping Configuration.
- Configure SMB. See SMB Configuration.
- Configure access control, ACL entries, and ACL inheritance on shares.

SMB and NFSv3 do not use the same access control model. For best results, configure the ACL on the root directory from a SMB client because the SMB access control model is a more verbose model. For information on inheritable trivial ACL entries, see Access Control Lists for Filesystems.

SMB DFS Namespaces

The Distributed File System (DFS) is a virtualization technology delivered over the SMB and MSRPC protocols. DFS allows administrators to group shared folders located on different servers by transparently connecting them to one or more DFS namespaces. A DFS namespace is a virtual view of shared folders in an organization. An administrator can select which shared folders to present in the namespace, design the hierarchy in which those folders appear and determine the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside in a single, high-capacity file system. Users can navigate the folders in the namespace without needing to know the server names or shared folders hosting the data.

Only one share per system may be provisioned as a standalone DFS namespace. Domain-based DFS namespaces are not supported. Note that one DFS namespace may be provisioned per cluster, even if each cluster node has a separate storage pool. To provision a SMB share as a DFS namespace, use the DFS Management MMC Snap-in to create a standalone namespace.

When Oracle ZFS Storage Appliance is not joined to an Active Directory domain, additional configuration is necessary to allow Workgroup users to modify DFS namespaces. To enable an SMB local user to create or delete a DFS namespace, that user must have a separate local account created on the server. For information about steps to let the SMB local user `dfsadmin` manipulate DFS namespaces, see Adding DFS Namespaces to a Local SMB Group.

SMB Microsoft Stand-alone DFS Namespace Management Tools Support Matrix

The following table lists operations (subcommands/options) of the Microsoft DFS tools on various Windows operating system versions. It identifies which of these are supported by the DFS service on Oracle ZFS Storage Appliance for managing a standalone DFS namespace on the appliance.

- **`y`** - supported
- **`n`** - not supported
- **`NA`** - not applicable

```
Microsoft Windows systems         XP|2003|2003|Vista|2008|2008|Win7|
                                    |    |  R2|     |    |  R2|    |
                                 SP3| SP2| SP2|  SP2| SP2| SP1| SP1|
                                    |    |    |     |    |    |    |
dfscmd CLI:                         |    |    |     |    |    |    |
                                    |    |    |     |    |    |    |
/map [comment] [/restore]          y|   y|   y|    y|   y|   y|   y|
/unmap                             y|   y|   y|    y|   y|   y|   y|
/add [/restore]                    y|   y|   y|    y|   y|   y|   y|
/remove                            y|   y|   y|    y|   y|   y|   y|
/view [/partial | /full]           y|   y|   y|    y|   y|   y|   y|
                                    |    |    |     |    |    |    |
                                    |    |    |     |    |    |    |
dfsutil CLI (old format):           |    |    |     |    |    |    |
                                    |    |    |     |    |    |    |
/addstdroot [/comment]             y|   y|   y|    n|   n|   y|   y|
/remstdroot                        y|   y|   y|    n|   n|   y|   y|
/root:<DfsName> /view              n|   n|   n|    y|   y|   y|   y|
/addlink [/comment]               NA|  NA|  NA|    y|   y|   y|   y|
/removelink                       NA|  NA|  NA|    y|   y|   y|   y|
/state /display                   NA|  NA|  NA|    y|   y|   y|   y|
/state /enable                    NA|  NA|  NA|    y|   y|   y|   y|
/state /disable                   NA|  NA|  NA|    y|   y|   y|   y|
/ttl /display                     NA|  NA|  NA|    y|   y|   y|   y|
/ttl /set                         NA|  NA|  NA|    y|   y|   y|   y|
/server:<MachineName> /view        y|   y|   y|    y|   y|   y|   y|
                                    |    |    |     |    |    |    |
                                    |    |    |     |    |    |    |
dfsutil CLI (new format):           |    |    |     |    |    |    |
                                    |    |    |     |    |    |    |
root addstd [comment]             NA|  NA|  NA|    n|   n|   y|   y|
root remove                       NA|  NA|  NA|    n|   n|   y|   y|
root (view namespace)             NA|  NA|  NA|    y|   y|   y|   y|
link add [comment]                NA|  NA|  NA|    y|   y|   y|   y|
link remove                       NA|  NA|  NA|    y|   y|   y|   y|
link (view)                       NA|  NA|  NA|    y|   y|   y|   y|
target add                        NA|  NA|  NA|    y|   y|   y|   y|
target remove                     NA|  NA|  NA|    y|   y|   y|   y|
target (view)                     NA|  NA|  NA|    y|   y|   y|   y|
property comment (view)           NA|  NA|  NA|    y|   y|   y|   y|
property comment set              NA|  NA|  NA|    y|   y|   y|   y|
property ttl (view)               NA|  NA|  NA|    y|   y|   y|   y|
property ttl set                  NA|  NA|  NA|    y|   y|   y|   y|
property state (view)             NA|  NA|  NA|    y|   y|   y|   y|
property state offline            NA|  NA|  NA|    y|   y|   y|   y|
property state online             NA|  NA|  NA|    y|   y|   y|   y|
                                    |    |    |     |    |    |    |
                                    |    |    |     |    |    |    |
DFS GUI:                            |    |    |     |    |    |    |
                                    |    |    |     |    |    |    |
add standalone root                y|   y|   y|    n|   n|   n|   n|
remove standalone root             y|   y|   y|    n|   n|   n|   n|
change root comment                y|   y|   y|    n|   n|   n|   n|
change root timeout                y|   y|   y|    n|   n|   n|   n|
add link                           y|   y|   y|    n|   n|   n|   n|
remove link                        y|   y|   y|    n|   n|   n|   n|
change link comment                y|   y|   y|    n|   n|   n|   n|
change link timeout                y|   y|   y|    n|   n|   n|   n|
add link's target                  y|   y|   y|    n|   n|   n|   n|
remove link's target               y|   y|   y|    n|   n|   n|   n|
enable link's referral (target)    y|   y|   y|    n|   n|   n|   n|
disable link's referral (target)   y|   y|   y|    n|   n|   n|   n|
hide root                          y|   y|   y|    y|   y|   y|   y|
show root                          y|   y|   y|    y|   y|   y|   y|
display links                      y|   y|   y|    n|   n|   n|   n|
display targets                    y|   y|   y|    n|   n|   n|   n|
                                  XP|2003|2003|Vista|2008|2008|Win7|
                                    |    |  R2|     |    |  R2|    |
                                 SP3| SP2| SP2|  SP2| SP2| SP1| SP1|
```

Note that:

- Oracle Solaris does not verify the DFS link target.
- CLI commands for modifying and viewing comment and timeout (TTL) are applicable to both root and link.
- CLI commands for viewing state are applicable to root, root's target, link, and link's target.
- CLI commands for modifying state are only applicable for link and link's target.

Adding DFS Namespaces to a Local SMB Group

Use the following procedure to add a DFS namespace to a local SMB group.

1. Create a local user account on the server for user `dfsadmin`. Be sure to use the same password as when the local user was first created on the Windows machine.
2. Add `dfsadmin` to local SMB group Administrators.
3. Log in as `dfsadmin` on the Windows machine from which the DFS namespace will be modified.

SMB Autohome

For Windows file sharing, Autohome provides access to filesystems using the SMB protocol. Autohome defines and maintains home directory shares for users that access the system through SMB. Autohome rules map SMB clients to home directories.

Setting Autohome Rules:

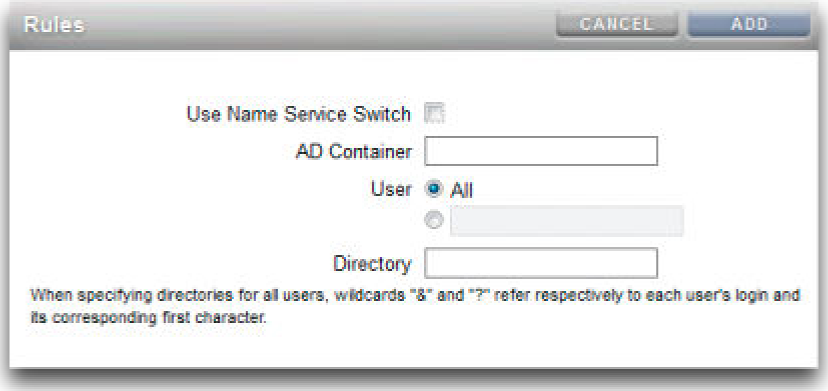

Use Name Service Switch □

AD Container

User ● All

○

Directory

When specifying directories for all users, wildcards "&" and "?" refer respectively to each user's login and its corresponding first character.

- **Use Name Service Switch -** Toggles Name Service Switch (NSS) on or off. You cannot create an NSS rule and an rule for all users at the same time.
- **AD Container -** Sets the Active Directory container, for example: `dc=com,dc=fishworks,ou=Engineering,CN=myhome`.

- **User** - Sets the Autohome rule for all All users or for the user you specify. When you specify a user, the wildcards "&" and "?" refer to a user's login and its corresponding first character.
- **Directory** - Sets the directory for the rule, for example: /export/wdp.

Adding SMB Autohome Rules (CLI)

Use the following procedure to add SMB autohome rules.

1. Go to configuration services smb.
2. Use the create command to add autohome rules, and the list command to list existing rules.

   This example adds a rule for the user Chris then lists the rules:

```
hostname:> configuration services smb
hostname:configuration services smb> create
hostname:configuration services rule (uncommitted)> set use_nss=false
hostname:configuration services rule (uncommitted)> set user=Chris
hostname:configuration services rule (uncommitted)> set directory=/export/wdp
hostname:configuration services rule (uncommitted)> set
container="dc=com,dc=fishworks,
   ou=Engineering,CN=myhome"
hostname:configuration services rule (uncommitted)> commit
hostname:configuration services smb> list
RULE       NSS      USER         DIRECTORY            CONTAINER
rule-000   false    Chris        /export/wdp          dc=com,dc=fishworks,
   ou=Engineering,CN=myhome
```

3. Create autohome rules using wildcard characters.

   The & character matches the users' username, and the ? character matches the first letter of the users' username. The following example uses wildcards to match all users:

```
hostname:configuration services smb> create
hostname:configuration services rule (uncommitted)> set use_nss=false
hostname:configuration services rule (uncommitted)> set user=*
hostname:configuration services rule (uncommitted)> set directory=/export/?/&
hostname:configuration services rule (uncommitted)> set
container="dc=com,dc=fishworks,
   ou=Engineering,CN=myhome"
hostname:configuration services rule (uncommitted)> commit
hostname:configuration services smb> list
RULE       NSS      USER         DIRECTORY            CONTAINER
rule-000   false    Chris        /export/wdp          dc=com,dc=fishworks,
   ou=Engineering,CN=myhome
```

4. The name service switch can also be used to create autohome rules:

```
hostname:configuration services smb> create
hostname:configuration services rule (uncommitted)> set use_nss=true
hostname:configuration services rule (uncommitted)> set
container="dc=com,dc=fishworks,
   ou=Engineering,CN=myhome"
hostname:configuration services rule (uncommitted)> commit
hostname:configuration services smb> list
RULE       NSS      USER         DIRECTORY            CONTAINER
rule-000   true                                       dc=com,dc=fishworks,
   ou=Engineering,CN=myhome
```

Adding a User to an SMB Local Group

Local groups are groups of domain and/or local users that grant additional privileges to those users.

SMB Local Groups:

- **Administrators** - Administrators can bypass file permissions to change the ownership on files.
- **Backup Operators** - Backup Operators can bypass file access controls to backup and restore files.

1. Go to `configuration services smb groups`.

   ```
   hostname:configuration services smb> groups
   ```

2. Enter `create`.

   ```
   hostname:configuration services smb groups> create
   ```

3. Specify the user you want to add to the group.

   ```
   hostname:configuration services smb member (uncommitted)> set user=Chris
   ```

4. Enter the group name, and then commit the change.

   ```
   hostname:configuration services smb member (uncommitted)> set group="Backup
   Operators"
   hostname:configuration services smb member (uncommitted)> commit
   ```

5. Enter `list` to confirm the user was added to the specified group.

   ```
   hostname:configuration services smb groups> list
   MEMBER       USER                         GROUP
   member-000   WINDOMAIN\Chris              Backup Operators
   ```

SMB MMC Integration

The Microsoft Management Console (MMC) is an extensible framework of registered components, known as snap-ins, that provide comprehensive management features for both the local system and remote systems on the network. Computer Management is a collection of Microsoft Management Console tools, that may be used to configure, monitor and manage local and remote services and resources.

To use the MMC functionality on Oracle ZFS Storage Appliance in workgroup mode, be sure to add the Windows administrator who will use the management console to the Administrators local group on the appliance. Otherwise you may receive an `Access is denied` or similar error on the administration client when attempting to connect to the appliance using the MMC.

The appliance supports the following Computer Management facilities: The Event Viewer MMC snap-in displays the Application log, Security log, and System log. These logs show the contents of the alert, audit, and system logs of the appliance.

The following screen shows an example of the Application log and the properties dialog box for an error event.

SMB Event Viewer:

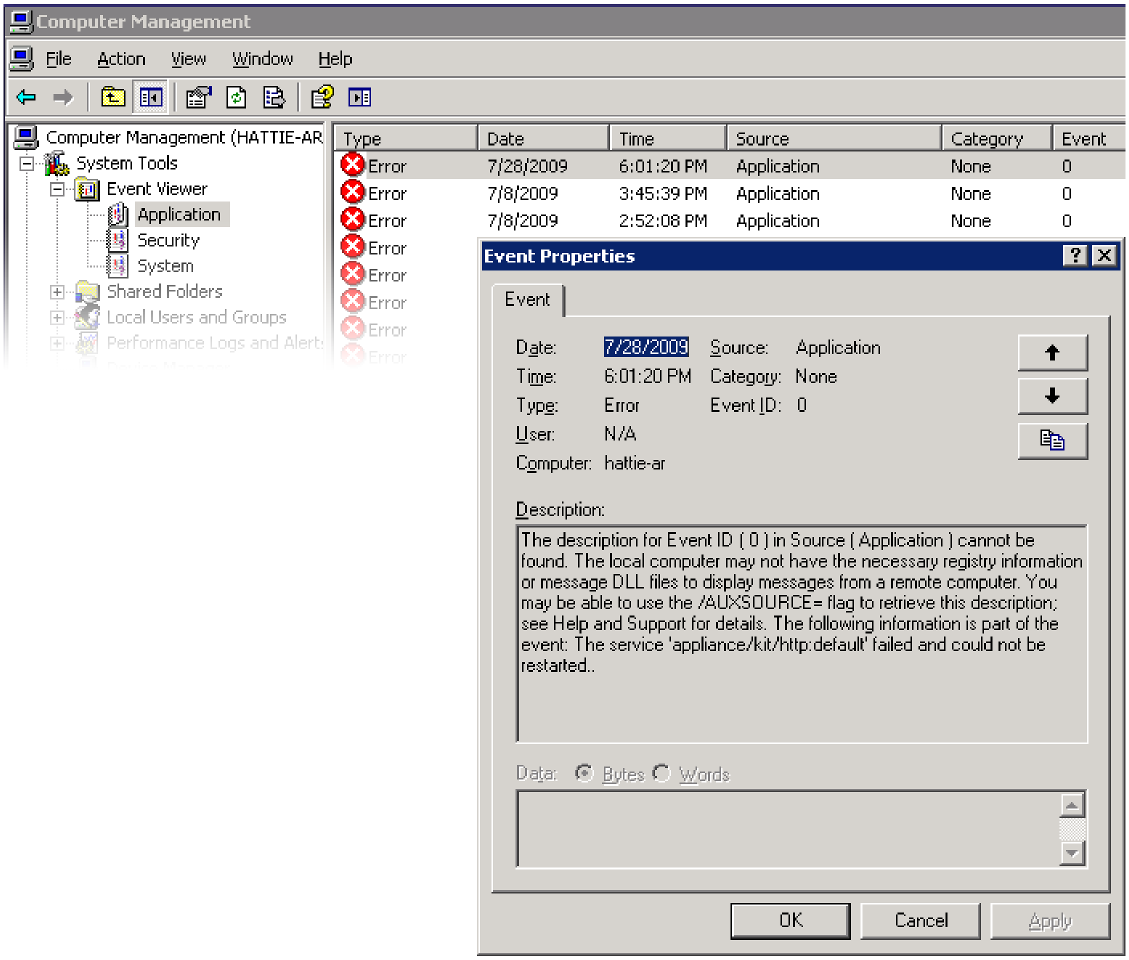

Type	Date	Time	Source	Category	Event
Error	7/28/2009	6:01:20 PM	Application	None	0
Error	7/8/2009	3:45:39 PM	Application	None	0
Error	7/8/2009	2:52:08 PM	Application	None	0
Error					
Error					
Error					
Error					
Error					

Date: 7/28/2009 Source: Application
Time: 6:01:20 PM Category: None
Type: Error Event ID: 0
User: N/A
Computer: hattie-ar

Description:
The description for Event ID (0) in Source (Application) cannot be found. The local computer may not have the necessary registry information or message DLL files to display messages from a remote computer. You may be able to use the /AUXSOURCE= flag to retrieve this description; see Help and Support for details. The following information is part of the event: The service 'appliance/kit/http:default' failed and could not be restarted..

Data: ● Bytes ○ Words

SMB Share Management

Support for SMB share management includes the following:

- Listing shares
- Setting ACLs on shares
- Changing share permissions
- Setting the description of a share

Features not currently supported via MMC include the following:

- Adding or deleting a share
- Setting client side caching property
- Setting maximum allowed or number of users property

The following screen shows an example of permission properties for a share:

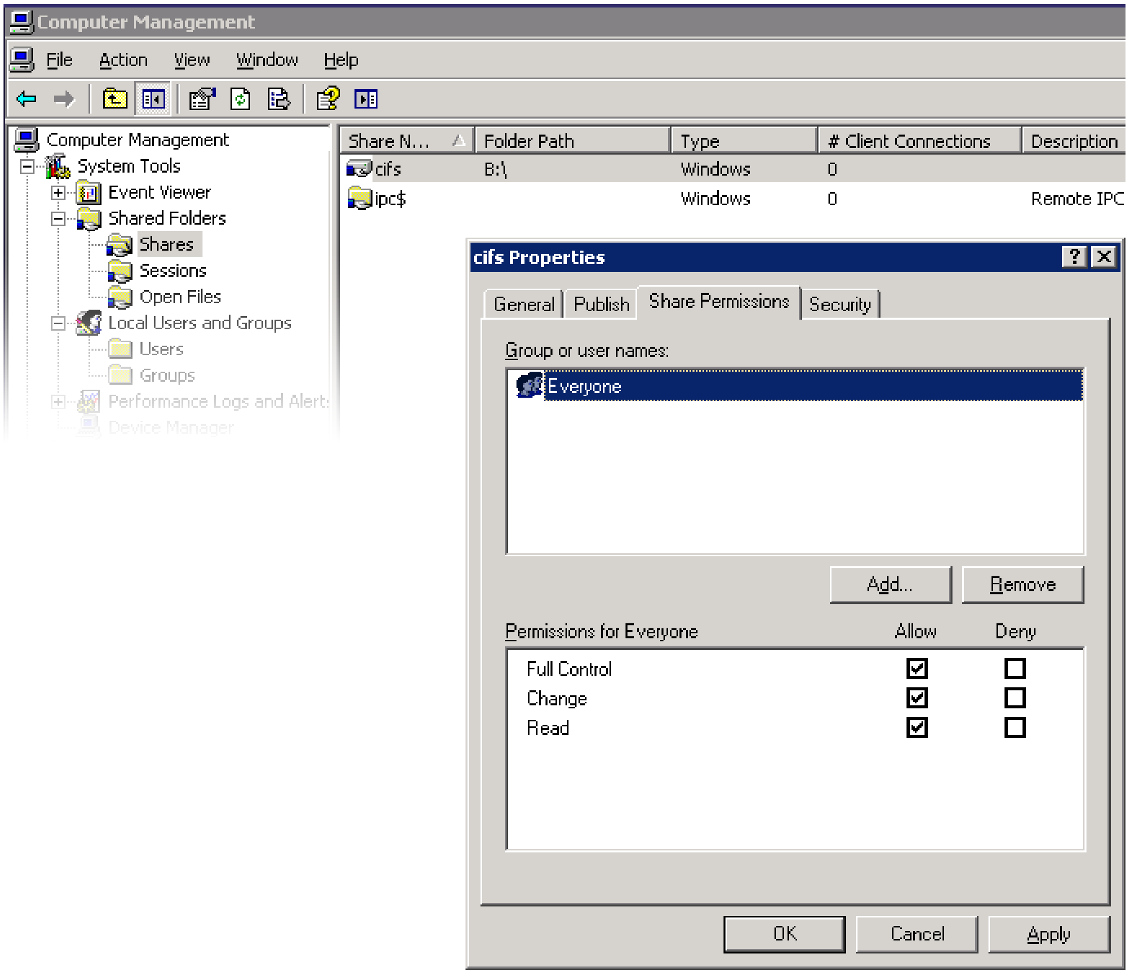

Share N...	Folder Path	Type	# Client Connections	Description
cifs	B:\	Windows	0	
ipc$		Windows	0	Remote IPC

Group or user names:

Everyone

Permissions for Everyone	Allow	Deny
Full Control	☑	□
Change	☑	□
Read	☑	□

SMB Users, Groups, and Connections

The following SMB features are supported:

- Viewing local SMB users and groups
- Listing user connections, including listing the number of open files per connection
- Closing user connections
- Listing open files, including listing the number of locks on the file and file open mode
- Closing open files

The following screen shows an example of open files per connection:

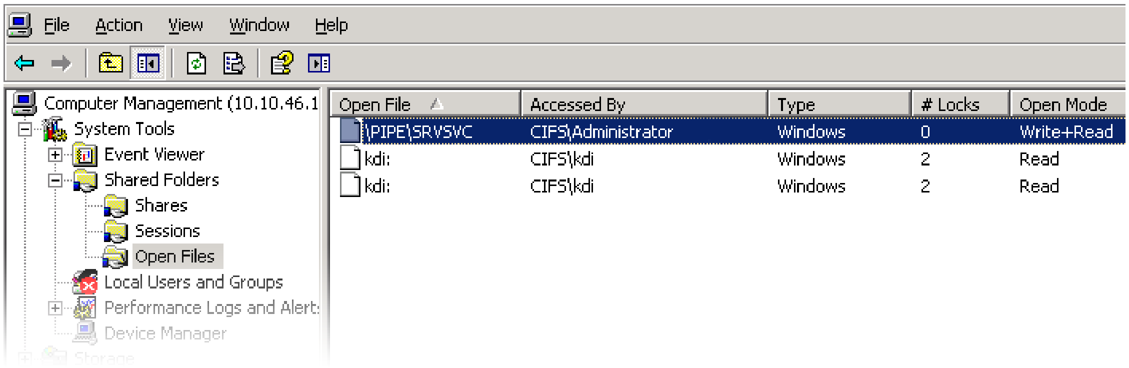

Open File	Accessed By	Type	# Locks	Open Mode
\PIPE\SRVSVC	CIFS\Administrator	Windows	0	Write+Read
kdi:	CIFS\kdi	Windows	2	Read
kdi:	CIFS\kdi	Windows	2	Read

The following screen shows an example of open sessions:

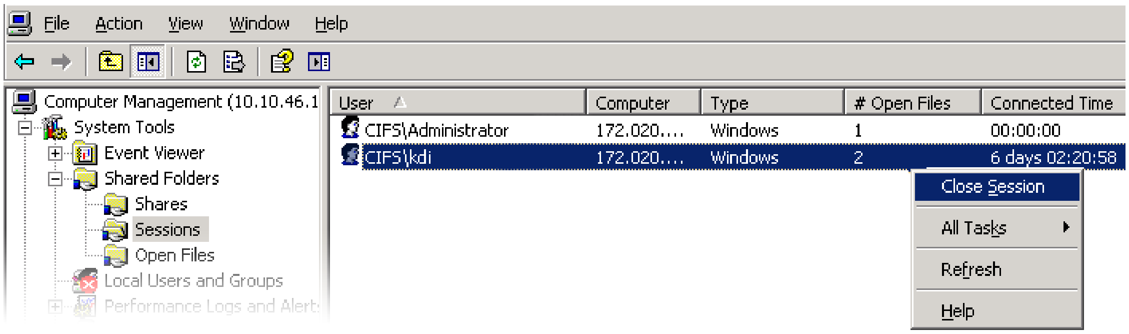

User	Computer	Type	# Open Files	Connected Time
CIFS\Administrator	172.020....	Windows	1	00:00:00
CIFS\kdi	172.020....	Windows	2	6 days 02:20:58

Listing SMB Services

Listing of appliance services is supported using the MMC application. However, you cannot enable or disable services using the MMC application. Support includes listing of appliance services. Services cannot be enabled or disabled using the Computer Management MMC application.

The following screen shows an example of general properties for the `vscan` service:

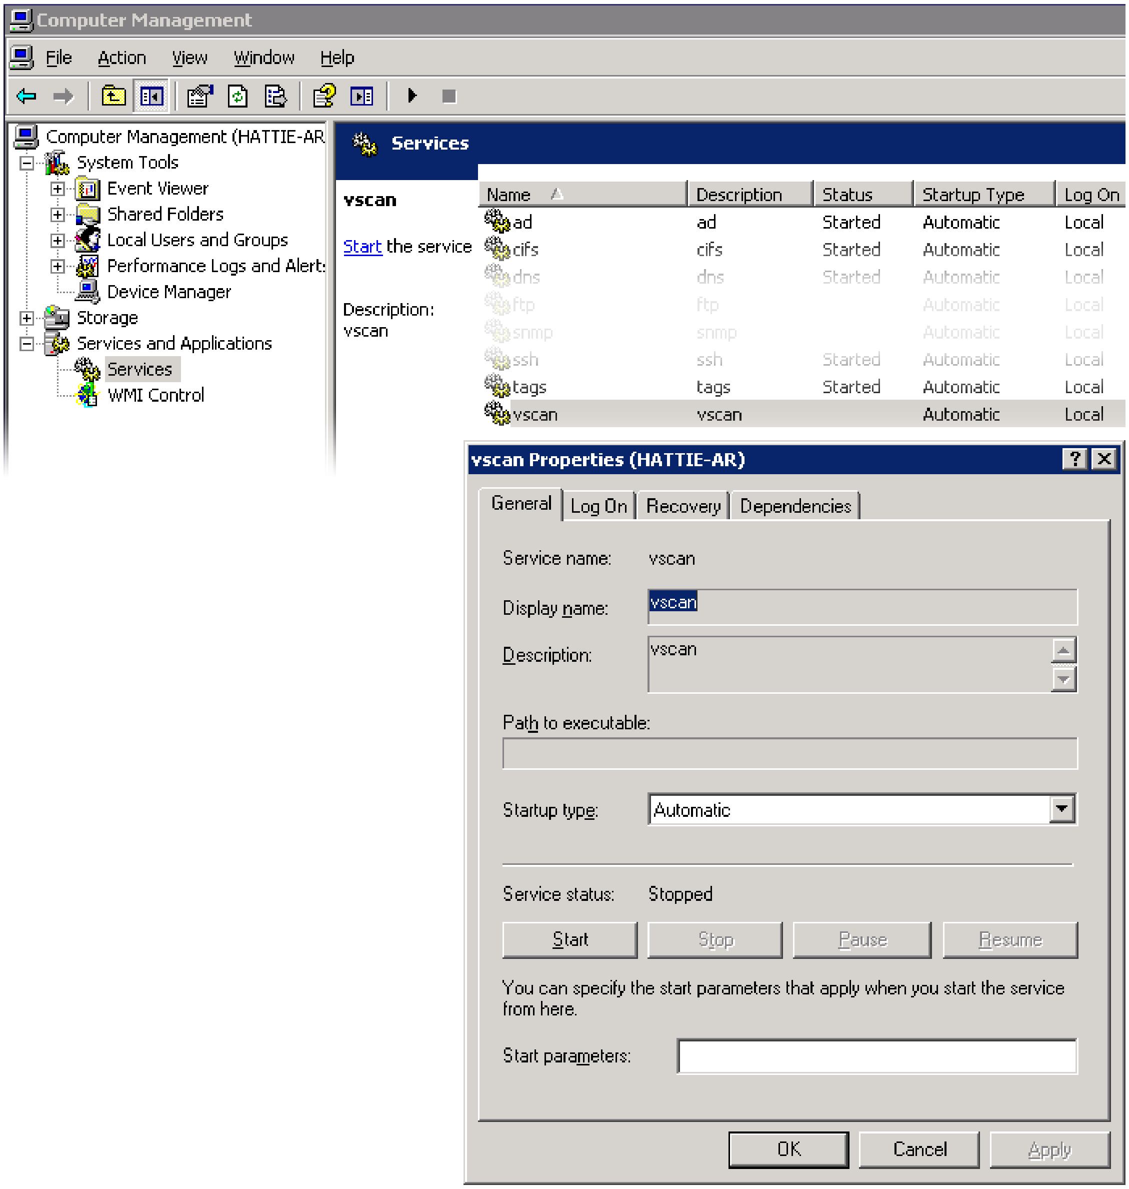

Name	Description	Status	Startup Type	Log On
ad	ad	Started	Automatic	Local
cifs	cifs	Started	Automatic	Local
dns	dns	Started	Automatic	Local
ftp	ftp		Automatic	Local
snmp	snmp		Automatic	Local
ssh	ssh	Started	Automatic	Local
tags	tags	Started	Automatic	Local
vscan	vscan		Automatic	Local

vscan Properties (HATTIE-AR)

Service name: vscan

Display name: vscan

Description: vscan

Path to executable:

Startup type: Automatic

Service status: Stopped

You can specify the start parameters that apply when you start the service from here.

Start parameters:

To ensure that only the appropriate users have access to administrative operations, there are some access restrictions on the operations performed remotely using MMC.

Table 3-45   Users and Allowed Operations

User	Allowed Operations
Regular users	List shares
Members of the Administrators or Power Users groups	Manage shares, list user connections
Members of the Administrators group	List open files and close files, disconnect user connections, view services and event log

Configuring SMB (BUI)

Initial configuration of Oracle ZFS Storage Appliance may be completed using the BUI or the CLI, and should take less than 20 minutes. Initial setup may also be performed again later using the **Maintenance: System** contexts of the BUI or CLI. Initial configuration takes you through the following steps.

1. **Configure Network Devices, Datalinks, and Interfaces.**
   1. Create interfaces using the **Datalink** add or **Interface** add icon ⊕ or by using drag-and-drop of devices to the **Datalink** or **Interface** list.
   2. Set the desired properties, and click **Apply** to add them to the list.
   3. Set each interface to **active** or **standby** as appropriate.
   4. Click **APPLY** at the top of the page to commit your changes.
2. **Configure DNS.**
   1. Provide the base domain name.
   2. Provide the IP address of at least one server that is able to resolve hostname and server records in the Active Directory portion of the domain namespace.
3. **Configure NTP authentication keys to ensure clock synchronization.**
   1. Click the add icon ⊕ to add a new key.
   2. Specify the number, type, and private value for the new key and apply the changes. The key appears as an option next to each specified NTP server.
   3. Associate the key with the appropriate NTP server and apply the changes. To ensure clock synchronization, configure the appliance and the SMB clients to use the same NTP server.
4. **Specify Active Directory as the directory service for users and groups.**
   1. Set the directory domain.
   2. Click **APPLY** to commit your changes.
5. **Configure a storage pool.**
   1. Click the add icon ⊕ to add a new pool.
   2. Set the pool name.
   3. On the **Allocate and verify storage** screen, configure the disk shelf allocation for the storage pool. If no disk shelves are detected, check your disk shelf cabling and power.
   4. Click **COMMIT** to advance to the next screen.
   5. On the **Configure Added Storage** screen, select the desired data profile. Each is rated in terms of availability, performance and capacity. Use these ratings to determine the best configuration for your business needs.
   6. Click **COMMIT** to activate the configuration.
6. **Configure Remote Support.**
   1. If the appliance is not directly connected to the Internet, configure an HTTP proxy through which the remote support service may communicate with Oracle.
   2. Enter your Online Account user name and password. A privacy statement will be displayed for your review.

  c. Choose one of your inventory teams with which to register. The default team for each account is the same as the account user name, prefixed with a `$` symbol.

  d. Commit your initial configuration changes.

Configuring SMB Active Directory (BUI)

Use the following procedure to configure SMB Active Directory for Oracle ZFS Storage Appliance.

1. Create an account for the appliance in the Active Directory domain.

   For detailed instructions, refer to Active Directory Configuration.

2. From the **Configuration** menu, select **Services**, then **Active Directory**, and click **Join Domain**.
3. Specify the Active Directory domain, administrative user, and administrative password.
4. Click **APPLY** to commit the changes.

Configuring SMB Project and Share (BUI)

Use the following procedure to configure a SMB project and share.

1. From the **Shares** menu, select **Shares**.
2. Create a project.

   a. On the **Shares** screen, click the panel open icon to expand the **Projects** panel.

   b. Click the add icon ⊕ to add a new project.

   c. Specify the project name, and click **APPLY**.

3. Select the new project from the **Projects** panel.
4. Click the add item icon ⊕ to add a filesystem.
5. Click the edit icon ✎ for the filesystem.
6. Click the **General** tab, and deselect the **Inherit from project** check box.
7. Choose a mountpoint under **/export**, even though SMB shares are accessed by resource name, and click **APPLY**.
8. Click the **Protocols** tab for the project, and set the SMB resource name to **on**.
9. Enable **sharesmb** and **share-level ACL** for the project.
10. Click **APPLY** to activate the configuration.

Configuring SMB Data Service (BUI)

Use the following procedure to configure the SMB data service.

1. From the **Configuration** menu, select **Services**, then **SMB**, and click the power icon ⏻ to enable the service.
2. Set SMB properties, and click **APPLY** to activate the configuration. See SMB Service Properties.

3. Click the **Autohome** tab on the **Configuration: Services: SMB** screen to set autohome rules to map SMB clients to home directories, as described in SMB Autohome. Click **APPLY** to activate the configuration.
4. Click the **Local Groups** tab on the **Configuration: Services: SMB** screen, and use the add item icon ⊕ to add administrators or backup operator users to local groups, as described in Adding a User to an SMB Local Group. Click **APPLY** to activate the configuration.

SMTP Configuration

The SMTP service sends all mail generated by the appliance, typically in response to alerts as configured on the **Alerts** screen. The SMTP service does not accept external mail; it only sends mail generated automatically by the appliance itself.

By default, the SMTP service uses DNS (MX records) to determine where to send mail. If DNS is not configured for the appliance's domain, or the destination domain for outgoing mail does not have DNS MX records setup properly, the appliance can be configured to forward all mail through an outgoing mail server, commonly called a smarthost.

Table 3-46 SMTP Properties

Property	Description
Send mail through smarthost	If enabled, all mail is sent through the specified outgoing mail server. Otherwise, DNS is used to determine where to send mail for a particular domain.
Smarthost hostname	Outgoing mail server hostname.
Allow customized from address	If enabled, the `From` address for email is set to the `Custom from` address property. It might be desirable to customize this if the default `From` address is being identified as spam, for example.
Custom from address	The `From` address to use for outbound email.
SMTP Authentication	Defines the Authentication service extension. Options: • `No Authentication` (default; `ANONYMOUS` in the CLI) • `Plain`
Auth Username	The authorized username if `Plain` authentication is enabled.
Auth Password	The authorized user's password if `Plain` authentication is enabled.

When changing properties, you can use Alerts to send a test email to verify that the properties are correct. A common reason for undelivered email is misconfigured DNS, which prevents the appliance from determining which mail server to deliver the mail to. As described previously, a smarthost could be used if DNS cannot be configured.

Table 3-47 SMTP Logs

Log	Description
`network-smtp:sendmail`	Logs the SMTP service events.
`mail`	Log of SMTP activity (including mails sent).

SNMP Configuration

The Simple Network Management Protocol (SNMP) service provides two different functions on Oracle ZFS Storage Appliance:

- Appliance status information can be served by SNMP.
- Alerts can be configured to send SNMP traps. See Configuring Alerts.

SNMP versions v1, v2c, and v3 are available when this service is enabled. The appliance supports a maximum of 128 physical and logical network interfaces. More than 128 network interfaces could cause time outs for such commands as `snmpwalk` and `snmpget`. If you need more than 128 network interfaces, contact Oracle Support.

To configure SNMP, see the following sections:

- Configuring SNMP to Serve Appliance Status (BUI)
- Configuring SNMP to Send Traps (BUI)
- SNMP Properties
- SNMP MIBs
- Sun FM MIB
- Sun AK MIB

Configuring SNMP to Serve Appliance Status (BUI)

Use the following procedure to configure SNMP to serve the Oracle ZFS Storage Appliance status.

1. From the **Configuration** menu, select **Services**, then **SNMP**.
2. Set the **Community name**, **Authorized network**, and **Contact string**.
3. Optional: Set the **Trap destination(s)** to a remote SNMP host, else set this to `127.0.0.1`.
4. Click **APPLY** to commit the configuration.

Configuring SNMP to Send Traps (BUI)

Use the following procedure to configure SNMP to send traps.

1. From the **Configuration** menu, select **Services**, then **SNMP**.
2. Set the **Community name**, **Contact string**, and **Trap destination(s)**.
3. Optional: Set the **Authorized network** to allow SNMP clients, else set this to `127.0.0.1/8`.
4. Click **APPLY** to commit the configuration.
5. You must configure alerts to send the traps you want to receive.

   For more information about alerts, see Configuring Alerts.

Related Topics

SNMP Properties

SNMP Properties

The following table describes the SNMP properties.

Table 3-48 SNMP Properties

Property	Description
Version	Toggles between `v1/2c` and `v3`.
Community name	Toggles between `public` and `user-input`. If you select `user-input`, you must also enter a `community name`. If you select `v3`, this property is not available.
Authorized network/subnet	Enter an appropriate IPv4 address and subnet (integers from 0-32). If you select `v3`, this property is not available.
Appliance contact	Enter an appropriate appliance contact.
Username/password	Enter a valid username (max 501 characters) and password (8-501 characters). If you select `v1/2c`, this property is not available.
Authentication	Toggles between `MD5` and `SHA` authentication algorithms. If you select `v1/2c`, this property is not available. Note that `MD5` is no longer recommended and might be removed in a future release.
Privacy	Toggles among `None`, `DES`, and `AES` encryption algorithms. If you select `v1/2c`, this property is not available. Note that `DES` is no longer recommended and might be removed in a future release.
Engine ID	The `EngineID` value hashed by `snmpd`. If SNMP was not previously enabled, the label shows `0x000`.
Trap destinations	Lets you add IPv4 addresses. Use the “+” and “-” buttons to add or remove addresses.

The SNMP service also provides the MIB-II location string. This property is sourced from the System Identity configuration.

SNMP MIBs

If the SNMP services is online, authorized networks will have access to the following MIBs (Management Information Bases):

Table 3-49 SNMP MIBs

MIB	Purpose
`.1.3.6.1.2.1.1`	MIB-II system - generic system information, including hostname, contact and location
`.1.3.6.1.2.1.2`	MIB-II interfaces - network interface statistics
`.1.3.6.1.2.1.4`	MIB-II IP - Internet Protocol information, including IP addresses and route table
`.1.3.6.1.2.1.31`	MIB-II ifMIB - extensions to the generic interfaces structure
`.1.3.6.1.4.1.42`	Sun Enterprise MIB (SUN-MIB.mib.txt)
`.1.3.6.1.4.1.42.2.195`	Sun FM - fault management statistics (SUN-FM-MIB.mib.txt)
`.1.3.6.1.4.1.42.2.225`	Sun AK - appliance information and statistics (SUN-AK-MIB.mib.txt)

> **Note:**
>
> Sun SNMP MIB files are available at `https://`*`your IP address or host name`*`:215/help/docs/snmp/SUN-MIB.mib.txt`.

Sun FM MIB

The Sun FM MIB (`SUN-FM-MIB.mib`) provides access to SUN Fault Manager information such as:

- Active problems on the system
- Fault Manager events
- Fault Manager configuration information

There are four main tables to read:

Table 3-50 Sun FM MIBs

OID	Contents
`.1.3.6.1.4.1.42.2.195.1.1`	Fault Management problems
`.1.3.6.1.4.1.42.2.195.1.2`	Fault Management fault events
`.1.3.6.1.4.1.42.2.195.1.3`	Fault Management module configuration
`.1.3.6.1.4.1.42.2.195.1.5`	Fault Management faulty resources

See the MIB file for the full descriptions.

> **Note:**
>
> Sun FM MIB files are available at `https://`*`your IP address or host name`*`:215/help/docs/snmp/SUN-FM-MIB.mib.txt`.

Sun AK MIB

The Sun AK MIB (`SUN-AK-MIB.mib`) provides the following information:

- Product description string and part number
- Appliance software version
- Appliance and chassis serial numbers
- Install, update and boot times
- Cluster state, including peer node
- Share status for both filesystems and LUNs (pool name, project name, share name, size, used and available gigabytes and bytes, filesystem mountpoint)
- Replica share status (pool name, project name, share name, size, used and available bytes, replica share's source name, filesystem mountpoint)
- Pool status (name, profile, status, total size, available size, used size by type, data compression and data deduplication ratios)

- Hardware status for disks (component name, faulted, present, enclosing chassis name, vendor, model, serial number, speed, type)

There are six main tables to read:

Table 3-51 Sun AK MIBs

OID	Contents
`.1.3.6.1.4.1.42.2.225.1.4`	General appliance information
`.1.3.6.1.4.1.42.2.225.1.5`	Cluster status
`.1.3.6.1.4.1.42.2.225.1.6`	Share status
`.1.3.6.1.4.1.42.2.225.1.7`	Replica share status
`.1.3.6.1.4.1.42.2.225.1.8`	Pool status
`.1.3.6.1.4.1.42.2.225.1.9`	Hardware status

See the MIB file for the full descriptions.

Note:

Sun AK MIB files are available at `https://`*`your IP address or host name`*`:215/help/docs/snmp/SUN-AK-MIB.mib.txt`.

SRP Configuration

When you configure a LUN on Oracle ZFS Storage Appliance, you can export that volume over a SCSI Remote Protocol (SRP) target. The SRP service allows initiators to access targets using the SRP protocol.

For information on SRP targets and initiators, see Configuring Storage Area Network (SAN).

SSH Configuration

The SSH (Secure Shell) service allows users to log in to the Oracle ZFS Storage Appliance CLI and perform most of the same administrative actions that can be performed in the BUI. The SSH service can also be used as a way to execute automated scripts from a remote host, such as for retrieving daily logs or Analytics statistics.

SSH keys can be configured for individual accounts using the preferences function described in Setting Appliance Preferences.

SSH can also be used in conjunction with Kerberos authentication. For information about the appliance Kerberos service, see Kerberos Configuration.

For added security when configuring SSH, you can specify the ciphers and MACs, as described in SSH Properties and Logs.

To configure SSH, see the following sections:

- Disabling root SSH Access (CLI)
- SSH Properties and Logs

Disabling root SSH Access (CLI)

Use the following procedure to disable root SSH access on Oracle ZFS Storage Appliance.

1. Go to `configuration services ssh`.
2. Set `permit root login` to `false`.
3. Commit the configuration.

SSH Properties and Logs

The following tables describe the SSH properties and logs.

Table 3-52 SSH Properties

Property	Description	Examples
Login grace period	The SSH connection will be disconnected after this many seconds if the client has failed to authenticate	`120`
Permit root login	Allows the root user to log in using SSH	`yes`
Port (for incoming connections)	The designated port for incoming connections	`22`

Table 3-53 SSH Security Properties

Property	Description
Ciphers	Ciphers for SSH connections.
MACs	Message authentication codes (MACs) for SSH connections.

The SSH service events log is available in `network-ssh:default`.

Syslog Configuration

The Syslog Relay service provides two different functions on Oracle ZFS Storage Appliance:

- Alerts can be configured to send Syslog messages to one or more remote systems. See Configuring Alerts.
- Services on the appliance that are syslog capable will have their syslog messages forwarded to remote systems.

A *syslog message* is a small event message transmitted from the appliance to one or more remote systems (or as we like to call it: intercontinental `printf`). The message contains the following elements:

- A facility describing the type of system component that emitted the message.
- A severity describing the severity of the condition associated with the message.
- A timestamp describing the time of the associated event in UTC.
- A hostname describing the canonical name of the appliance.
- A tag describing the name of the system component that emitted the message. See SYSLOG Alert Message Format for details of the message format.

- A message describing the event itself. See SYSLOG Alert Message Format for details of the message format.

Syslog receivers are provided with most operating systems, including Oracle Solaris and Linux. A number of third-party and open-source management software packages also support Syslog. Syslog receivers allow administrators to aggregate messages from a number of systems on to a single management system and incorporated into a single set of log files.

The Syslog Relay can be configured to use the "classic" output format described by RFC 3164, or the newer, versioned output format described by RFC 5424. Syslog messages are transmitted as UDP datagrams. Therefore they are subject to being dropped by the network, or may not be sent at all if the sending system is low on memory or the network is sufficiently congested. Administrators should therefore assume that in complex failure scenarios in a network some messages may be missing and were dropped.

Syslog Properties

- **Protocol Version** (`version`) - The version of the Syslog protocol to use, either Classic Syslog (RFC 3164) or Updated Syslog (RFC 5424).
- **Destinations** (`dst`) - The list of destination IPv4, IPv6, and FQDN addresses to which messages are relayed.
- **Audit Classes** (`audit_classes`) - Controls which audit events are relayed to the remote target. The default setting does not specify an option, which are as follows:
  - **Administrative Audit** (`AdministrativeAudit`) - Oracle ZFS Storage Appliance audit events.
  - **Per File Audit** (`PerFileAudit`) - Per-file audit for Oracle Solaris audit events.
  - **Login/Logout Audit** (`LoginLogoutAudit`) - Log in and log out Oracle Solaris audit events.

To configure syslog, see the following sections:

- Classic Syslog: RFC 3164
- Updated Syslog: RFC 5424
- SYSLOG Message Format
- SYSLOG Alert Message Format
- Example Configuring an Oracle Solaris Receiver (CLI)
- Example Configuring a Linux Receiver (CLI)

Classic Syslog: RFC 3164

The Classic Syslog protocol includes the facility and level values encoded as a single integer priority, the timestamp, a hostname, a tag, and the message body.

The tag will be one of the tags described in SYSLOG Message Format.

The hostname will be the canonical name of the appliance as defined by the System Identity configuration. For more information, see System Identity Configuration.

Updated Syslog: RFC 5424

The Classic Syslog protocol includes the facility and level values encoded as a single integer priority, a version field (1), the timestamp, a hostname, a app-name, and the message body.

Syslog messages relayed by the storage systems will set the RFC 5424 procid, msgid, and structured-data fields to the nil value (-) to indicate that these fields do not contain any data.

The app-name will be one of the tags described in SYSLOG Message Format.

The hostname will be the canonical name of the appliance as defined by the System Identity configuration. For more information, see System Identity Configuration.

SYSLOG Message Format

The Syslog protocol itself does not define the format of the message payload, leaving it up to the sender to include any kind of structured data or unstructured human-readable string that is appropriate. Oracle ZFS Storage Appliance systems use the syslog subsystem tag ak to indicate a structured, parseable message payload, described next. Other subsystem tags indicate arbitrary human-readable text, but administrators should consider these string forms *unstable* and subject to change without notice or removal in future releases of the storage software.

Table 3-54 SYSLOG Message Formats

Facility	Tag Name	Description
daemon	`ak`	Generic tag for appliance subsystems. All alerts will be tagged `ak`, indicating a `SUNW-MSG-ID` follows.
daemon	`idmap`	Identity Mapping service for POSIX and Windows identity conversion. See Identity Mapping Configuration.
daemon	`smbd`	SMA Data Protocol for accessing shares. See SMB Configuration.

SYSLOG Alert Message Format

If an alert is configured with the Send Syslog Message action, it will produce a syslog message payload containing localized text consisting of the following standard fields. Each field will be prefixed with the field name in CAPITAL letters followed by a colon and white (negative) space character.

Table 3-55 SYSLOG Alert Message Formats

Field Name	Description
`SUNW-MSG-ID`	The stable Sun Fault Message Identifier associated with the alert. Each system condition and fault diagnosis that produces an administrator alert is assigned a persistent, unique identifier in Sun's Fault Message catalog. These identifiers can be easily read over the phone or scribbled down in your notebook, and link to a corresponding knowledge article found at My Oracle Support (https://support.oracle.com/) "Predictive Self-Healing" (Doc ID 1154428.1).
`TYPE`	The type of condition. This will be one of the labels: `Fault`, indicating a hardware component or connector failure; `Defect` indicating a software defect or misconfiguration; `Alert`, indicating a condition not associated with a fault or defect, such as the completion of a backup activity or remote replication.
`VER`	The version of this encoding format itself. This description corresponds to version "1" of the `SUNW-MSG-ID` format. If a "1" is present in the `VER` field, parsing code may assume that all of the subsequent fields will be present. Parsing code should be written to handle or ignore additional fields if a decimal integer greater than one is specified.

Table 3-55 (Cont.) SYSLOG Alert Message Formats

Field Name	Description
SEVERITY	The severity of the condition associated with the problem that triggered the alert. The list of severities is shown in the next table.
EVENT-TIME	The time corresponding to this event. The time will be in the form Day Mon DD HH:MM:SS YYYY in UTC. For example: Wed Aug 17 21:34:22 2022.
PLATFORM	The platform identifier for the appliance. This field is for Oracle Service use only.
CSN	The chassis serial number of the appliance.
HOSTNAME	The canonical name of the appliance as defined by the System Identity configuration. See System Identity.
SOURCE	The subsystem within the appliance software that emitted the event. This field is for Oracle Service use only.
REV	The internal revision of the subsystem. This field is for Oracle Service use only.
EVENT-ID	The Universally Unique Identifier (UUID) associated with this event. Oracle's Fault Management system associates a UUID with each alert and fault diagnosis such that administrators can gather and correlated multiple messages associated with a single condition, and detect duplicate messages. Oracle Service personnel can use the EVENT-ID to retrieve additional postmortem information associated with the problem that may help Oracle respond to the issue.
DESC	Description of the condition associated with the event.
AUTO-RESPONSE	The automated response to the problem, if any, by the Fault Management software included in the system. Automated responses include capabilities such as proactively offlining faulty disks, DRAM memory chips, and processor cores.
REC-ACTION	The recommended service action. This will include a brief summary of the recommended action, but administrators should consult the knowledge article and this documentation for information on the complete repair procedure.

The SEVERITY field will be set to one of the following values:

Table 3-56 SYSLOG Severity Fields

Severity	Syslog Level	Description
Minor	LOG_WARNING	A condition occurred that does not currently impair service, but the condition needs to be corrected before it becomes more severe.
Major	LOG_ERR	A condition occurred that does impair service but not seriously.
Critical	LOG_CRIT	A condition occurred that seriously impairs service and requires immediate correction.

Example Configuring an Oracle Solaris Receiver (CLI)

Most operating systems include a syslog receiver, but some configuration steps may be required to turn it on. Consult the documentation for your operating system or management software for specific details of syslog receiver configuration.

Oracle Solaris includes a bundled syslogd that can act as a syslog receiver, but the remote receive capability is disabled by default. To enable Oracle Solaris to receive syslog traffic, use svccfg and svcadm to modify the syslog settings as follows:

```
# svccfg -s system/system-log setprop config/log_from_remote = true
# svcadm restart system/system-log
```

The Oracle Solaris syslogd only understands the classic Syslog protocol. Refer to the Oracle Solaris syslog.conf(4) man page for information on how to configure filtering and logging of the received messages.

By default, Oracle Solaris syslogd records messages to /var/adm/messages and a test alert would be recorded as follows:

```
Aug 14 21:34:22 poptart.example.us.com poptart ak: SUNW-MSG-ID: AK-8000-LM, \
TYPE: alert, VER: 1, SEVERITY: Minor\nEVENT-TIME: Wed Aug 14 21:34:22 2019\n\
PLATFORM: i86pc, CSN: 12345678, HOSTNAME: poptart\n\
SOURCE: jsui.359, REV: 1.0\n\
EVENT-ID: 92dfeb39-6e15-e2d5-a7d9-dc3e221becea\n\
DESC: A test alert has been posted.\n\
AUTO-RESPONSE: None.\nIMPACT: None.\nREC-ACTION: None.
```

Example Configuring a Linux Receiver (CLI)

Most operating systems include a syslog receiver, but some configuration steps may be required to turn it on. Consult the documentation for your operating system or management software for specific details of syslog receiver configuration.

Most Linux distributions include a bundled sysklogd(8) daemon that can act as a syslog receiver, but the remote receive capability is disabled by default. To enable Linux to receive syslog traffic, edit the /etc/sysconfig/syslog configuration file such that the -r option is included (enables remote logging):

```
SYSLOGD_OPTIONS="-r -m 0"
```

and then restart the logging service:

```
# /etc/init.d/syslog stop
# /etc/init.d/syslog start
```

Some Linux distributions have an ipfilter packet filter that will reject syslog UDP packets by default, and the filter must be modified to permit them. On these distributions, use a command similar to the following to add an INPUT rule to accept syslog UDP packets:

```
# iptables -I INPUT 1 -p udp --sport 514 --dport 514 -j ACCEPT
```

By default, Linux syslogd records messages to /var/log/messages and a test alert would be recorded as follows:

```
Aug 12 22:03:15 192.168.1.105 poptart ak: SUNW-MSG-ID: AK-8000-LM, \
TYPE: alert, VER: 1, SEVERITY: Minor EVENT-TIME: Mon Aug 12 22:03:14 2019 \
PLATFORM: i86pc, CSN: 12345678, HOSTNAME: poptart SOURCE: jsui.3775, REV: 1.0 \
EVENT-ID: 9d40db07-8078-4b21-e64e-86e5cac90912 \
DESC: A test alert has been posted. AUTO-RESPONSE: None. IMPACT: None. \
REC-ACTION: None.
```

System Identity Configuration

This service provides configuration for the system name, location, and an optional system login message that is displayed upon log in. You might need to change these if Oracle ZFS Storage Appliance is moved to a different network location, or repurposed. You can change this data in the BUI by going to **Configuration: Service: System Identity**. To access the same data in the CLI, go to the `configuration service identity` context.

System Identity Properties and Logs

The System Identity properties are described in the following table.

Table 3-57 System Identity Properties

BUI Text	CLI Property	Description
System Name	`nodename`	A single canonical identifying name for the appliance that is shown in the user interface. This name is separate from any DNS names that are used to connect to the system (which would be configured on remote DNS servers). This name can be changed at any time.
System Location	`syslocation`	A text string to describe where the appliance is physically located. If SNMP is enabled, this will be exported as the `syslocation` string in MIB-II.
System Login Message	`loginmessage`	The optional system login message is configurable at any time. The message is displayed in both the BUI and the CLI. In the BUI, the message displays on the upper left of the initial login screen. Users must click the **Continue** button before they can log in. In the CLI, the message displays immediately above the prompt after login.

Changing services properties is documented in Setting Service Properties (BUI) and Setting Service Properties (CLI). The CLI property names are shorter versions of those listed above.

The System Identity service events log is available in `system-identity:node`.

To view service logs, refer to Using Logs in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

TFTP Configuration

Trivial File Transfer Protocol (TFTP) is a simple protocol to transfer files. TFTP is designed to be small and easy to implement; therefore, it lacks most of the features of a regular FTP. TFTP only reads and writes files (or mail) from/to a remote server. It cannot list directories, and currently has no provisions for user authentication.

Table 3-58 TFTP Properties

Property	Description
Default Root Directory	The TFTP login location. The default is `/export` and points to the top of the shares hierarchy. All users will be logged into this location after successfully authenticating with the TFTP service.

To use TFTP with a share, see Adding TFTP Access to a Share (BUI).

Adding TFTP Access to a Share (BUI)

Use the following procedure to add TFTP access to a share.

1. From the **Configuration** menu, select **Services**.
2. Check that the TFTP service is enabled and online. If not, enable the service.
3. From the **Shares** menu, select **Shares**, and select or add a share.
4. Go to the **Protocols** tab, and check that TFTP access is enabled.
5. Optional: Set the **Share mode access** to **Read only** or **Read/Write**.

Virus Scan Configuration

The Virus Scan service scans for viruses at the filesystem level. When a file is accessed from any protocol, the Virus Scan service first scans the file, and both deny access and quarantine the file if a virus is found. After a file has been scanned with the latest virus definitions, it is not rescanned until it is next modified. Files accessed by NFS clients that have cached file data, or been delegated read privileges by the NFSv4.0 or NFSv4.1 servers, may not be immediately quarantined.

To configure Virus Scan, see the following sections:

- Configuring Virus Scanning for a Share (BUI)
- Virus Scan Properties and Logs
- Virus Scan File Extensions
- Scanning Engines

Configuring Virus Scanning for a Share (BUI)

Use the following procedure to configure virus scanning for a share.

1. From the **Configuration** menu, select **Services**, then **Virus Scan**.
2. Enable the service.
3. Set the appropriate properties.
4. Click **APPLY** to commit the configuration.
5. From the **Shares** menu, edit a filesystem or a project.
6. Select the **General** tab, and enable the virus scan option.

Related Topics

- Virus Scan Properties and Logs
- Virus Scan File Extensions
- Scanning Engines

Virus Scan Properties and Logs

This section describes the virus scan properties and logs.

Table 3-59 Virus Scan Properties

Property	Description
Maximum file size to scan (maxsize)	Files larger than this size will not be scanned, to avoid significant performance penalties. These large files are unlikely to be executable themselves (such as database files), and so are less likely to pose a risk to vulnerable clients. The default value is 1 GB.
Allow access to files that exceed maximum file size (maxsize_action)	Enabled by default, this property allows access to files larger than the maximum scan size (which are therefore not scanned prior to being returned to clients). Administrators at a site with more stringent security requirements may elect to disable this option and increase the maximum file size, so that all accessible files are known to be scanned for viruses.
Use TLS (use_tls)	Enabled by default, this property determines if the TLS protocol is used to connect to virus scan engines, which must be TLS enabled. If non-TLS scan engines were set up on the appliance before software release OS8.8.69 (when this feature was supported), this property is disabled after a software update to OS8.8.69 or later. To use TLS virus scan engines, enable this property, upload TLS certificates, and assign the certificates as trusted certificates to the virus scan (**vscan**) service. Also, see Scanning Engines.
Require that certificate matches hostname (host_match)	Enabled by default, this property determines if the TLS trusted certificate must match the hostname of the virus scan engine server.

The Virus Scan service events log is vscan.

Related Topics

- Configuring Virus Scanning for a Share (BUI)
- Virus Scan File Extensions
- Scanning Engines

Virus Scan File Extensions

This section describes how to control which files are scanned. The default value, *, causes all files to be scanned. Because scanning all files may impact performance, you can designate a subset of files to scan.

For example, to scan only high-risk files, including zip files, but not files with names that match the pattern data-archive*.zip, you could configure the settings shown in the following table.

Table 3-60 Settings to Scan Only Specified High Risk File Types

Action	File Extension Pattern
Scan	exe
Scan	com
Scan	bat
Scan	doc
Scan	zip

Table 3-60 (Cont.) Settings to Scan Only Specified High Risk File Types

Action	File Extension Pattern
`Don't Scan`	data-archive*.zip
`Don't Scan`	*

Use `Don't Scan *` to exclude all other file types not explicitly included in the scan list. A file named `file.name.exe.bat.jpg123` would *not* be scanned, because only the `jpg123` portion of the name, the extension, would be compared against the rules.

Do *not* use exclude settings before include settings. For example, do not use a `Don't Scan *` setting before `Scan` settings since that would exclude all file types that come after it.

Related Topics

- Configuring Virus Scanning for a Share (BUI)
- Virus Scan Properties and Logs
- Scanning Engines

Scanning Engines

In this section, specify which scanning engines to use. A scanning engine is an external third-party virus scanning server that Oracle ZFS Storage Appliance contacts to scan files using either the TLS or ICAP (Internet Content Adaptation Protocol, RFC 3507) protocol.

To communicate with a TLS-enabled server, upload one or more TLS trusted certificates and assign them to the virus scan (**vscan**) service, as described in Configuring Certificates.

Table 3-61 Scanning Engines Properties

Property	Description
Enable	Use this scan engine.
Host	Hostname or IP address of the scan engine server.
Maximum Connections	Maximum number of concurrent connections. Some scan engines operate better with connections limited to 8.
Port	Port for the scan engine.

Related Topics

- Configuring Virus Scanning for a Share (BUI)
- Virus Scan Properties and Logs
- Virus Scan File Extensions

4
Shares and Projects

Oracle ZFS Storage Appliance uses storage pools, projects, and shares to organize data. Shares are filesystems and LUNs that are exported over supported data protocols to clients of the appliance. All shares within a project can share common settings, and quotas can be enforced at the project level in addition to the share level. For more information about how the appliance organizes data, see About Storage Pools, Projects, and Shares.

To create and modify projects, use these tasks:

- Creating a Project - BUI, CLI
- Editing a Project - BUI, CLI
- Renaming a Project - BUI, CLI
- Deleting a Project - BUI, CLI

To create and modify filesystems and LUNs, use these tasks:

- Creating a Filesystem or LUN in a Project - BUI, CLI
- Editing a Filesystem or LUN - BUI, CLI
- Renaming a Filesystem or LUN - BUI, CLI
- Moving a Filesystem or LUN to a Different Project - BUI, CLI
- Deleting a Filesystem or LUN - BUI, CLI
- Setting User or Group Quotas - BUI, CLI

To understand more about how the appliance organizes storage, see these topics:

- About Storage Pools, Projects, and Shares
- Space Management for Shares
- Project and Share Properties
- File Retention Management
- Working with Filesystem Namespace
- Share Usage Statistics
- Share and Project Protocols
- Access Control Lists for Filesystems
- Working with Schemas
- Snapshots and Clones
- Remote Replication

Creating a Project (BUI)

A project inherits properties of the parent pool. A share inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A project that is created in an encrypted pool is automatically encrypted. You cannot create an unencrypted project in an encrypted pool.

Use this procedure to create an unencrypted project. To create an encrypted project in either an unencrypted pool or an encrypted pool, see Creating an Encrypted Project (BUI).

1. From the **Shares** menu, select **Projects**.
2. Click the add icon ⊕ next to **Projects** or in the expanded **Projects** panel. To expand the **Projects** panel, click the arrow icon .
3. In the **Create Project** dialog box, enter a name for the new project.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters **_ - . :**
4. Set properties as appropriate for this project.

   Project properties are described in Project Properties.

   > ⚠ **Caution:**
   >
   > If setting file retention, first review section File Retention Management, which includes the required user role authorizations for certain features. Plan for the future: Actively retained files cannot be modified, even after their retention has expired. Also, mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.
5. Click **APPLY**.

   The new project is added to the **Projects** list.

Related Topics

- Project Properties
- Creating an Encrypted Project (BUI)
- File Retention Management

Creating a Project (CLI)

A project inherits properties of the parent pool. A share inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A project that is created in an encrypted pool is automatically encrypted. You cannot create an unencrypted project in an encrypted pool.

Use this procedure to create an unencrypted project. To create an encrypted project in either an unencrypted pool or an encrypted pool, see Creating an Encrypted Project (CLI).

1. Go to `shares`.

   ```
   hostname:> shares
   ```
2. Enter the `project` command and a name for the project.

The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **- . :**

```
hostname:shares> project home
```

3. Use the `get` and `set` commands to set properties as appropriate for this project.

   Project properties are described in Project Properties.

   > **Caution:**
   >
   > If setting file retention, first review section File Retention Management, which includes the required user role authorizations for certain features. Plan for the future: Actively retained files cannot be modified, even after their retention has expired. Also, mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.

4. Enter `commit`.

```
hostname:shares home> commit
```

Related Topics

- Project Properties
- Creating an Encrypted Project (CLI)
- File Retention Management

Editing a Project (BUI)

To modify project properties, use these steps.

1. From the **Shares** menu, select **Projects**.
2. Select a project in one of the following ways:
   - Hover over the project and click the edit icon ✎ .
   - Double-click the project name.
   - Click the arrow icon next to **Projects** to expand the panel, then click on the project name.

   The project is selected, and tabs are displayed for editing the properties.
3. Click one of the tabs to edit the project properties.
4. Modify the project properties for the project as needed.

   Project properties are described in Project Properties.

   If file retention is set, review section File Retention Management. Retained projects and their files cannot be modified, even after their expiration. However, with the proper user role authorization, you can edit the file retention periods.

Related Topics

- Snapshots and Clones
- Remote Replication

- File Retention Management

Editing a Project (CLI)

To modify project properties, use these steps.

1. Go to `shares`.

   ```
   hostname:> shares
   ```

2. Enter `select` and a project name.

   ```
   hostname:shares> select home
   ```

3. To list the project properties, use the `get` command.
4. To modify the project properties, use the `set` command. Project properties and values are described in Project Properties. For example, to enable `vscan` for this project, enter the following command:

   ```
   hostname:shares home >set vscan=true
   ```

   If file retention is set, review section File Retention Management. Retained projects and their files cannot be modified, even after their expiration. However, with the proper user role authorization, you can edit the file retention periods.

5. Enter `commit`.

   ```
   hostname:shares home> commit
   ```

Related Topics

- Snapshots and Clones
- Remote Replication
- File Retention Management

Renaming a Project (BUI)

You can rename a project with file retention if the project is empty and has an unexpired grace period. If retention is set on an empty project and the grace period is met or not set, you cannot rename the project until the expiration period has been met. A non-empty project cannot be renamed until all retained filesystems within the project have an expired retention value. For more information, see File Retention Management.

> **Caution:**
>
> Changing a project name will disrupt active client I/O operations.

1. Disconnect any active clients connected to the project.
2. From the **Shares** menu, select **Projects**.
3. Click on the project name in the **Projects** list.
4. Enter the new name for the project.

The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters **_ - . :**

5. Press **Return**.
6. Click **OK** to confirm.

Related Topics

Project Properties

Renaming a Project (CLI)

You can rename a project with file retention if the project is empty and has an unexpired grace period. If retention is set on an empty project and the grace period is met or not set, you cannot rename the project until the expiration period has been met. A non-empty project cannot be renamed until all retained filesystems within the project have an expired retention value. For more information, see File Retention Management.

> **Caution:**
>
> Changing a project name will disrupt active client I/O operations.

1. Disconnect any active clients connected to the project.
2. Go to `shares`.

   ```
   hostname:> shares
   ```

3. To view the projects, use the `list` command.

   ```
   hostname:shares> list
   default
   home
   ```

4. Enter `rename`, the existing project name, and the new project name.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters **_ - . :**

   ```
   hostname:shares> rename home project1
   ```

5. To confirm the project was renamed, use the `list` command.

   ```
   hostname:shares> list
   default
   project1
   ```

Related Topics

Project Properties

Deleting a Project (BUI)

You can delete a project with file retention if the project is empty and has an unexpired grace period. If retention is set on an empty project and the grace period is met or not set, you cannot delete the project until the expiration period has been met. A non-empty project cannot be deleted until all retained filesystems within the project have an expired retention value. For more information, see File Retention Management.

> **Caution:**
>
> Deleting a project destroys all data in the project by deleting its filesystems and LUNs.

1. From the **Shares** menu, select **Projects**.
2. Hover over the project you want to delete, and click the destroy icon 🗑 .
3. Click **OK**.
4. To monitor the amount of space to be reclaimed in the storage pool if deferred update Asynchronous Dataset Deletion (OS8.7.0 or later) has been accepted, from the **Configuration** menu, select **Storage**, select the appropriate pool, and note the amount of space in field **Asynchronous Dataset Destroy**.

   When the operation has completed, **Asynchronous Dataset Destroy** is not displayed.

Deleting a Project (CLI)

You can delete a project with file retention if the project is empty and has an unexpired grace period. If retention is set on an empty project and the grace period is met or not set, you cannot delete the project until the expiration period has been met. A non-empty project cannot be deleted until all retained filesystems within the project have an expired retention value. For more information, see File Retention Management.

> **Caution:**
>
> Deleting a project destroys all data in the project by deleting its filesystems and LUNs.

1. Go to `shares`.

   ```
   hostname:> shares
   ```

2. Enter `destroy` and a project name.

   ```
   hostname:shares> destroy home
   ```

3. Enter `Y`.

   ```
   This will destroy all data in "home"! Are you sure? (Y/N)
   hostname:shares> Y
   ```

4. To monitor the amount of space to be reclaimed in the storage pool if deferred update Asynchronous Dataset Deletion (OS8.7.0 or later) has been accepted, enter `cd ..` to return to the root context. Enter `configuration storage` and enter `ls` to list storage pool properties. For the appropriate pool, note the amount of space for property `async_destroy_reclaim_space`.

   When the operation has completed, `0` (zero) is displayed.

Creating a Filesystem or LUN in a Project (BUI)

A filesystem or LUN that is created within a project inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A filesystem or LUN that is created within an encrypted project is automatically encrypted. You cannot create an unencrypted share in an encrypted project.

Use this procedure to create an unencrypted share. To create an encrypted share in either an unencrypted project or an encrypted project, see Creating an Encrypted Filesystem or LUN (BUI).

Before You Begin

If you are adding a filesystem or LUN to a non-default project, the project must already exist. To create a new project, see Creating a Project (BUI).

1. From the **Shares** menu, select **Shares**.
2. Select either **Filesystems** or **LUNs**.
3. Click the add icon ⊕ .
4. Complete the fields in the **Create Filesystem** or **Create LUN** dialog box.
   - For a filesystem, select a project and enter a name.
   - For a LUN, select a project, enter a name, and specify the volume size.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **- . :**
5. Set other properties.

   Share properties are described in Filesystem Properties and LUN Properties.

   > **Caution:**
   >
   > If setting file retention or if it is set at the project level, first review section File Retention Management, which includes the required user role authorizations for certain features. Plan for the future: Actively retained files cannot be modified, even after their retention has expired. Also, mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.

6. Click **APPLY**.

Related Topics

- Project and Share Properties
- Creating an Encrypted Filesystem or LUN (BUI)

Creating a Filesystem or LUN in a Project (CLI)

A filesystem or LUN that is created within a project inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A filesystem or LUN that is created within an encrypted project is automatically encrypted. You cannot create an unencrypted share in an encrypted project.

Use this procedure to create an unencrypted share. To create an encrypted share in either an unencrypted project or an encrypted project, see Creating an Encrypted Filesystem or LUN (CLI).

Before You Begin

If you are adding a filesystem or LUN to a non-default project, the project must already exist. To create a new project, see Creating a Project (CLI).

1. Go to `shares`.
2. Select the project.

   If the project that is selected is not the one you want, use the `select` *`project-name`* command to select a different project.
3. Create the filesystem or LUN.

   Enter `filesystem` *`filesystem-name`* or `lun` *`lun-name`*.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters **_ - . :**

   The following example creates a filesystem named `fs-1` in the `default` project.

   ```
   hostname:shares default> filesystem fs-1
   hostname:shares default/fs-1 (uncommitted)
   ```
4. If you are creating a LUN, enter `set volsize=` and the volume size.

   ```
   hostname:shares default/lun1 (uncommitted)> set volsize=2G
                                volsize = 2G (uncommitted)
   ```
5. Enter `commit`.

   ```
   hostname:shares default/fs-1 (uncommitted)> commit
   ```
6. Select the filesystem or LUN.

   Enter `select` *`filesystem-name`* or `select` *`lun-name`*.

   ```
   hostname:shares default> select fs-1
   ```
7. Use the `get` and `set` commands to set properties as appropriate for this filesystem or LUN.

   Share properties are described in Filesystem Properties and LUN Properties.

   > **Caution:**
   >
   > If setting file retention or if it is set at the project level, first review section File Retention Management, which includes the required user role authorizations for certain features. Plan for the future: Actively retained files cannot be modified, even after their retention has expired. Also, mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.
8. Enter `commit`.

Related Topics

- Project and Share Properties
- Creating an Encrypted Filesystem or LUN (CLI)

Editing a Filesystem or LUN (BUI)

To modify properties for an individual filesystem or LUN, use these steps.

1. From the **Shares** menu, select **Shares**.
2. Select **Filesystems** or **LUNs**.
3. Edit the filesystem or LUN.
   - Hover over the filesystem or LUN and click the edit icon ✎ .
   - Double-click the filesystem or LUN that you want to edit.

   The general properties are displayed for the filesystem or LUN.
4. Click the **Protocols**, **Access**, **Snapshots**, or **Replication** tab.
5. Modify the filesystem or LUN properties.

   Share properties are described in Filesystem Properties and LUN Properties.

   If file retention is set, review section File Retention Management. Retained files cannot be modified, even after their expiration. However, with the proper user role authorization, you can edit the file retention periods.
6. Click **APPLY**.

Related Topics

Project and Share Properties

Editing a Filesystem or LUN (CLI)

To modify properties for an individual filesystem or LUN, use these steps.

1. Select the filesystem or LUN that you want to change.
   1. Go to `shares`.
   2. Select the project.

      If the project that is selected is not the one you want, use the `select` *`project-name`* command to select a different project.
   3. Select the filesystem or LUN.

      Enter `select` *`filesystem-name`* or `select` *`lun-name`*.
2. Use the `get` command to list the properties of the share.
3. Use the `set` command to modify the filesystem or LUN properties.

   Share properties are described in Filesystem Properties and LUN Properties.

   If file retention is set, review section File Retention Management. Retained files cannot be modified, even after their expiration. However, with the proper user role authorization, you can edit the file retention periods.
4. Enter `commit`.

Related Topics

Project and Share Properties

Renaming a Filesystem or LUN (BUI)

Use the following procedure to rename a filesystem or LUN.

> **Caution:**
>
> Changing a share name will disrupt active client I/O operations.

1. Disconnect all active clients connected to the filesystem or LUN you want to rename.
2. From the **Shares** menu, select **Shares**.
3. Select **Filesystems** or **LUNs**.
4. Click on the filesystem or LUN name in the list.
5. Enter the new name for the filesystem or LUN.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **- . :**

   If file retention is set, review section File Retention Management. You cannot rename a retained file, even after its retention has expired.
6. Press **Return**.
7. Click **OK** to confirm.

Related Topics

- Filesystem Properties
- LUN Properties

Renaming a Filesystem or LUN (CLI)

Use the following procedure to rename a filesystem or LUN.

> **Caution:**
>
> Changing a share name will disrupt active client I/O operations.

1. Disconnect all active clients connected to the filesystem or LUN.
2. Go to `shares`.

   ```
   hostname:> shares
   ```
3. To view the projects, use the `list` command.

   ```
   hostname:shares> list
   default
   home
   ```

4. Enter `select` and the project name that contains the filesystem or LUN you want to rename.

   ```
   hostname:shares> select default
   ```

5. Enter `rename`, the existing filesystem or LUN name, and the new filesystem or LUN name.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (**.**) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **-.:**

   ```
   hostname:shares default> rename fs-1 fs-2
   ```

   If file retention is set, review section File Retention Management. You cannot rename a retained file, even after its retention has expired.

Related Topics

- Filesystem Properties
- LUN Properties

Moving a Filesystem or LUN to a Different Project (BUI)

Filesystems and LUNs within a project inherit the properties of the project.

If file retention is set, review section File Retention Management. You can move a retained filesystem to a different project. Note that mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.

1. From the **Shares** menu, select **Shares**.
2. Select **Filesystems** or **LUNs**.
3. Hover over the filesystem or LUN, and click the move icon ✥ .
4. Drag the filesystem or LUN to the different project under **Projects**.

   If the **Projects** panel is not expanded, the panel will automatically expand until the share is dropped onto a project.

Related Topics

- Filesystem Properties
- LUN Properties

Moving a Filesystem or LUN to a Different Project (CLI)

Filesystems and LUNs within a project inherit the properties of the project.

If file retention is set, review section File Retention Management. You can move a retained filesystem to a different project. Note that mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.

1. Go to `shares` and select the project that contains the filesystem or LUN to be moved.

   In this example, the `default` project contains the filesystem or LUN to be moved.

```
hostname> shares
hostname:shares> select default
```

2. Enter `move`, the name of the filesystem or LUN to be moved, and the name of the project to which to move it.

```
hostname:shares default> move foo home
```

Related Topics

- Filesystem Properties
- LUN Properties

Deleting a Filesystem or LUN (BUI)

If file retention is set, review section File Retention Management. A filesystem containing unexpired files with privileged file retention can be deleted. A filesystem containing unexpired files with mandatory file retention cannot be deleted until retention has expired for all of its retained files.

> **Caution:**
>
> Deleting a filesystem or LUN destroys all data in the share and cannot be undone.

1. From the **Shares** menu, select **Shares**.
2. Select **Filesystems** or **LUNs**.
3. Hover over the filesystem or LUN you want to delete, and click its destroy icon 🗑 .
4. Click **OK**.
5. To monitor the amount of space to be reclaimed in the storage pool if deferred update Asynchronous Dataset Deletion (OS8.7.0 or later) has been accepted, from the **Configuration** menu, select **Storage**, select the appropriate pool, and note the amount of space in field **Asynchronous Dataset Destroy**.

   When the operation has completed, **Asynchronous Dataset Destroy** is not displayed.

Related Topics

- Filesystem Properties
- LUN Properties

Deleting a Filesystem or LUN (CLI)

If file retention is set, review section File Retention Management. A filesystem containing unexpired files with privileged file retention can be deleted. A filesystem containing unexpired files with mandatory file retention cannot be deleted until retention has expired for all of its retained files.

> **Caution:**
>
> Deleting a filesystem or LUN destroys all data in the share and cannot be undone.

1. Go to `shares`.

   ```
   hostname> shares
   ```

2. Enter `select` and the project name that contains the filesystem or LUN.

   ```
   hostname:shares> select default
   ```

3. Enter `select` and the filesystem or LUN name.

   ```
   hostname:shares default>select fs-1
   ```

4. Enter `destroy`.

   ```
   hostname:shares default/fs-1> destroy
   This will destroy all data in "fs-1"! Are you sure? (Y/N)
   ```

5. Enter `Y`.

   ```
   hostname:shares default> Y
   ```

6. To monitor the amount of space to be reclaimed in the storage pool if deferred update Asynchronous Dataset Deletion (OS8.7.0 or later) has been accepted, enter `cd ../..` to return to the root context. Enter `configuration storage`, and enter `ls` to list storage pool properties. For the appropriate pool, note the amount of space for property `async_destroy_reclaim_space`.

   When the operation has completed, `0` (zero) is displayed.

Related Topics

- Filesystem Properties
- LUN Properties

Setting User or Group Quotas (BUI)

Quotas can be set for a user or group at the project or filesystem level.

1. From the **Shares** menu, select **Shares**, and select a project or share.
2. Click the **General** tab.
3. In the **Space Usage - Users & Groups** section, select **User**, **Group**, or **User or Group** from the drop-down menu.

   > **Note:**
   >
   > Any user that is not consuming any space on the filesystem, and does not have any quota set, does not appear in the list of active users.

4. To set a quota at the project level, select one of three options:
   - **None** - No quota is set for this filesystem.
   - **Default** - Sets the quota to the default quota at the project level; if no default was set, no quota is set for this filesystem.
   - Click the radio button, enter a quota in the **Size** field, and select a measurement.
5. Click **APPLY**.

   The user and group quota properties are validated separately from the other properties. However, you may only see one validation error if an invalid user/group as well as another invalid property is entered. Correcting one error and applying the changes will show any remaining error messages.

If you see an error message that an invalid property has been entered, it may be an invalid user/group, another invalid property on the page, or both. Fixing one invalid property and then applying the changes will reveal any remaining error messages.

Related Topics

Setting User or Group Quotas

Setting User or Group Quotas (CLI)

Quotas can be set for a user or group at the project or filesystem level.

1. Go to `shares`, select a project, then select a share, as shown in this example:

```
hostname:> shares select default select eschrock
```

2. Enter `users`, then `list` to see the current users.

```
hostname:shares default/eschrock> users
hostname:shares default/eschrock users> list
USER       NAME                            USAGE  QUOTA  SOURCE
user-000   root                             321K      -       -
user-001   ahl                             9.94K      -       -
user-002   eschrock                        20.0G      -       -
```

> **Note:**
>
> Any user that is not consuming any space on the filesystem, and does not have any quota set, does not appear in the list of active users.

3. Enter `select` and the `name=` of the user.

```
hostname:shares default/eschrock users> select name=eschrock
hostname:shares default/eschrock user-002> get
                          name = eschrock
                      unixname = eschrock
                        unixid = 132651
                       winname = (unset)
                         winid = (unset)
                         usage = 20.0G
                         quota = (unset)
                        source = (unset)
```

4. Enter `quota=` and a value. Enter `commit` and `done`.

> **Note:**
>
> To clear a quota, set the value to `0`.

```
hostname:shares default/eschrock user-002> set quota=100G
                         quota = 100G (uncommitted)
hostname:shares default/eschrock user-002> commit
hostname:shares default/eschrock user-002> done
```

5. To set a quota for such a user or group, use the quota command, after which the name and quota can be set.

   The SOURCE column displays local if the quota was set at the filesystem level, default if set at the project level, or - if no quota was set. In the following example, the default user quota set at the project level is 50 GB.

   If a default user or group quota was set at the project level, this procedure overrides that value.

```
hostname:shares default/eschrock users> quota
hostname:shares default/eschrock users quota (uncommitted)> set name=bmc
                          name = bmc (uncommitted)
hostname:shares default/eschrock users quota (uncommitted)> set quota=default
                         quota = default (uncommitted)
hostname:shares default/eschrock users quota (uncommitted)> commit
hostname:shares default/eschrock users> list
USER        NAME                             USAGE  QUOTA  SOURCE
user-000    root                              321K      -       -
user-001    ahl                              9.94K      -       -
user-002    eschrock                         20.0G   100G   local
user-003    bmc                                  -    50G  default
```

Related Topics

Setting User or Group Quotas

About Storage Pools, Projects, and Shares

Oracle ZFS Storage Appliance manages physical storage using a pooled storage model where all filesystems and LUNs share common space. This topic describes how storage is organized using storage pools, projects, and shares.

Storage Pools

Oracle ZFS Storage Appliance is based on the ZFS filesystem, which groups underlying storage devices into pools. Filesystems and LUNs, collectively referred to as shares, allocate from this storage pool as needed. Before creating filesystems or LUNs, you must first configure storage on the appliance. Once a storage pool is configured, there is no need to statically size filesystems, although this behavior can be achieved by using quotas and reservations.

While multiple storage pools are supported, this type of configuration is generally discouraged because it provides significant drawbacks as described in the Configuring Storage section. Multiple pools should only be used where the performance or reliability characteristics of two different profiles are drastically different, such as a mirrored pool for databases and a RAID-Z pool for streaming workloads.

When multiple pools are active on a single host, the BUI displays a drop-down list in the menu bar that can be used to switch between pools. In the CLI, the name of the current pool is displayed in parenthesis, and can be changed by setting the pool property. If only a single pool is configured, then these controls are hidden. When multiple pools are selected, the default pool chosen by the UI is arbitrary, so any scripted operation should be sure to set the pool name explicitly before manipulating any shares.

Projects

All filesystems and LUNs are grouped into projects. A project can be considered a consistency group that defines a common administrative control point for managing shares. All shares within a project can share common settings, and quotas can be enforced at the project level as well as at the share level. Projects can also be used solely for grouping logically related shares together, so their common attributes (such as accumulated space) can be accessed from a single point.

By default, the appliance creates a single default project when a storage pool is first configured. It is possible to create all shares within this default project, although for reasonably sized environments creating additional projects is strongly recommended, if only for organizational purposes.

Shares

Shares are filesystems and LUNs that are exported over supported data protocols to clients of the appliance. Exported filesystems can be accessed over SMB, NFS, HTTP/WebDav, and FTP. LUNs export block-based volumes and can be accessed over iSCSI or Fibre Channel.

The project/share is a unique identifier for a share within a pool. Projects within a pool cannot contain shares with the same name. If you attempt to name or rename a share using a name that is already in use, a mount point error occurs.

In addition to the default properties, you can configure shares and projects with any number of additional properties. These properties are given basic types for validation purposes, and are inherited like most other standard properties. The values are never consumed by the software in any way, and exist solely for end-user consumption. The property schema is global to the system, across all pools, and is synchronized between cluster peers.

Related Topics

- Space Management for Shares
- Project and Share Properties
- Snapshots and Clones

File Retention Management

The share file retention policy provides a facility for data governance, legal holds, and compliance records retention. Both data governance and regulatory compliance can be used to help protect from cyber and ransomware attacks.

- **Data Governance** - Data governance locks datasets (snapshots, objects or files) for a period of time, thus protecting the data from deletion. You might need to protect certain datasets as part of internal business process requirements or to protect datasets as part of your cyber-protection strategy. Data governance allows for adjustments in the retention strategy from privileged users.

  File retention data governance is implemented by creating a new project and filesystem with the "privileged" file retention policy. Privileged mode allows you to create a default retention setting for all new files, and to change that setting in the future to a shorter or longer duration. Files inherit the retention setting in effect when they are created. Retention can also be adjusted manually to a longer duration by changing the unlock timestamp. Projects and filesystems cannot be deleted when they have locked files.

- **Legal Holds** - A legal hold preserves certain business data in response to potential or ongoing lawsuits. A legal hold does not have a defined retention period, and it remains in effect until removed. Once the legal hold is removed, all protected data is immediately eligible for deletion unless other retention rules still apply.

  File retention legal holds on files are implemented by manually increasing the retention period on individual files, or by setting a hold for individual files so that their expiry date extends indefinitely. Because a legal hold may be required for an indefinite period of time, it is recommended to periodically extend manual retention. Holds on files never expire; the hold must be explicitly turned off. These two methods allow file retention to expire after the need for the legal hold has passed.

- **Regulatory Compliance** - Your industry might require you to retain a certain class of data for a defined length of time. Your data retention regulations might also require that you lock the retention settings. Regulatory compliance only allows you to increase the retention time, if at all. Regulatory compliance is the most restrictive locking strategy, and it often does not allow anyone, even an administrator, to make changes affecting retention.

  File retention regulatory compliance is implemented by creating a new project and filesystem with the "mandatory (no override)" file retention policy. Mandatory mode does not allow you to decrease the file retention duration. However, retention can be adjusted manually to a longer duration by changing the unlock timestamp. Regulatory compliance uses the same mechanisms as data governance, but it is much more restrictive. The project and filesystem cannot be deleted when locked files exist, and the storage pool cannot be unconfigured when locked files exist within the pool. This mode also requires usage of an NTP server, and the root user is locked out of remote access.

File retention completes the trio of retention products for Oracle ZFS Storage Appliance: file retention, snapshot retention, and object storage retention.

When the file retention policy is enabled, files become retained when set to readonly. Each file has a retained-until-expiration timestamp. This expiration date is either explicitly set or calculated based on the retention policy set on the share. When automatic retention is enabled, a file that has not been modified for the grace period is automatically retained at the default period value. Automatically retained files can have a longer retained-until-expiration date by manually setting the value.

A retained file cannot be modified, even after its expiration; this includes its name and attributes. When the expired date has been reached, retained files can be deleted, but not modified. A retained file's expiration date can be lengthened, but never shortened.

File retention is set at share creation, and it can be set to off (default), privileged, or mandatory. After setting privileged or mandatory file retention, you define the retention periods: minimum, maximum, default, and optional grace period.

This section contains the following topics:

Privileged File Retention

The privileged policy provides an override mechanism to allow retained files to be deleted before the expiry date. This requires the share to be exported with root access for NFS, and the user on the NFS client must be root.

A file's retained-until-expiration date can be set either manually before the file is marked as readonly or calculated based on last modification. When automatic retention is enabled, a file that has not been modified for the grace period is automatically retained for the time indicated by the default period. Files that are automatically retained can have their retained-until-expiration date manually set longer.

Unlike mandatory file retention, privileged file retention does not protect the share, project nor storage pool in which the retained file is contained. Therefore, a share, project, or storage pool containing actively retained files can be destroyed (share and project) or unconfigured (storage pool) at any time.

Mandatory File Retention

In addition to privileged file retention restrictions, mandatory file retention has further limitations. The `retentionMandatory` user role authorization is required to create a file with the mandatory retention policy. No user, not even the root user, can modify or delete a retained file. Also, shares with retained but unexpired files cannot be deleted, and storage pools containing files with unexpired retention cannot be unconfigured. A share with mandatory retention can be destroyed after all retained files have expired, and a storage pool with only expired retained files can be unconfigured. Also, a mandatorily retained file protects its ancestors and clone descendants from destruction.

> **Caution:**
>
> Mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.

You can change the retention period forward into the future but, unlike privileged file retention, you cannot change the retention backward in time. Thus, for example, you cannot immediately delete a retained filesystem, project, or its storage pool, but must wait until all files' retention times have expired.

To further protect mandatory retention shares, the software cannot be rolled back to a previous version without mandatory file retention nor can a storage pool's disks be reused while retention is active. Furthermore, NTP cannot be disabled, and NTP always synchronizes the clock, regardless of the skew amount. You cannot use a striped storage pool profile with mandatory retention because it does not provide redundancy. Also, the appliance factory reset feature is disabled until all mandatory file retention expires.

Automatic File Retention

The automatic file retention feature can be used with both mandatory and privileged file retention. When the grace period is set to a non-zero value, the file is automatically retained for the default retention period when the file has not been modified for the grace period. Until the grace period expires, the file can be modified—including its name, attributes, and ACL—or deleted. If necessary, the grace period can be modified by a user with the `retentionAuto` authorization.

If a filesystem is written to, the share's, project's, and pool's expiration is set to now plus the grace period plus the default retention period to protect the share, project, and storage pool. For this reason, after automatic retention is in effect, the grace period can be decreased, but not increased or unset.

Mandatory with Automatic File Retention Guidelines

Extra caution should be observed when using both mandatory and automatic file retention together because filesystems might be preserved for longer than anticipated or consume more storage resources than originally planned. Remember that mandatory file retention affects the filesystem, its project, and its storage pool.

For mandatory with automatic file retention, the share, project, and storage pool are protected for the grace period plus the default period when automatic retention is first enabled, and then extended with each write to a file. The files in the share must have no writes for the grace period plus the default period before the share, project or storage pool can be deleted or unconfigured.

Retention Period Settings

The retention period settings restrict the time values for retaining a file, and they are used for both mandatory and privileged file retention. If a file is retained in a mandatory retention policy filesystem, the file's retention time also affects its share, project, and storage pool. The administrator must have the `retentionPeriods` role authorization.

If no value is entered for a period, it is assumed to be zero. When setting a value, also set the time measurement unit, such as seconds, minutes or years.

- **Minimum file retention period** - The minimum amount of time for file retention. The default value is 0 (zero), and the value must be less than 100 years.
- **Maximum file retention period** - The maximum amount of time for file retention, which must be less than 100 years. The default value is 5 years.
- **Default file retention period** - The default amount of time for which a file is retained if it is automatically retained, or retained manually without first changing the file's access time attribute. After the default period has been met, the file cannot be modified, but it can be deleted. The default value must be between the minimum and maximum retention periods, inclusive. The default value is 0 (zero).
- **Automatic file retention grace period** - The amount of time a file must remain unmodified before it is automatically retained. This is known as the automatic file retention grace period or automatic file retention. When the grace period expires, the file is automatically retained at the default retention period setting. The grace period is not constrained by either the minimum period nor the maximum period, and it can only be modified by a user with the `retentionAuto` authorization. However, after the grace period is enabled, it can only be shortened, never extended, for the share.

File Retention on Expiry Policy

Use the file retention on expiry policy to automatically delete files after expiration or to place a hold on all retained files. The default setting is `off`, which does not apply any behavior after a file's retention expires, and the file remains on the system.

When set to `delete`, retained files are automatically deleted after expiration. The default file retention period must be a non-zero value. Optionally, set the property `retention.period.deletegrace` to the amount of time to delay automatically deleting the files.

When set to `hold`, the file's expiry date is extended indefinitely, and the `retention.policy.onexpiry` property must be changed to `off` or `delete` to be able to delete the file. Therefore, the hold setting is especially beneficial for files with legal holds.

If the `retention.policy.onexpiry` property is set to `hold` on a filesystem or project with mandatory file retention, the filesystem and project cannot be deleted, and the storage pool cannot be unconfigured. When viewing a pool's statistics, the number of held filesystems is displayed, along with the number of filesystems with mandatory retention.

> **Caution:**
>
> Mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.

The `hold` setting does not affect filesystems within a replication package nor does it affect the storage pool containing such replication packages. However, if a replication package containing a held filesystem is cloned, the replication package with the cloned and held filesystem is also held. Additionally, held filesystems in replication packages are not included the storage pool's held filesystems statistic.

To use the file retention on expiry policy, apply deferred update File Retention on Expiry, which is available in software version OS8.8.63 or later. For information on applying deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

Systems with software releases earlier than OS8.8.63 and that have not accepted the deferred update cannot import storage pools that use this feature; also, datasets cannot be received with property `retention.policy.onexpiry` set to anything except `off`.

User authorization `retentionOnexpiry` is required to set property `retention.policy.onexpiry`. See Assign Authorizations to User Roles.

Prerequisites

Before you can use the file retention feature, apply the File Retention deferred update, and assign user role authorizations. To set the policy for behavior after individual file retention expiry, apply deferred update File Retention on Expiry, and assign the appropriate user role authorization. Also, if mandatory file retention will be used, appropriately configure the appliance. Optionally, you can configure the appliance to allow root users on other systems to delete unexpired files with privileged retention via NFS. Also optionally, you can configure the appliance to allow SMB users in the Administrators group to delete unexpired files with privileged retention.

Apply the File Retention Deferred Updates

To use the file retention feature, apply deferred update File Retention, which is available in software release OS8.8.45 or later.

To set the policy for behavior after individual file retention expiry, apply deferred update File Retention on Expiry, which is available in software release OS8.8.63 or later.

For information on applying deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

Assign Authorizations to User Roles

Per the retention type, assign authorizations, under the Projects and Shares scope, to the appropriate user roles. For information on assigning authorizations to user roles, see Editing Authorizations for a Role - BUI, CLI.

These are the user role authorizations for file retention:

- `retentionMandatory` - User can enable mandatory file retention when creating a filesystem. The root user automatically has this authorization. After the appliance is configured for mandatory file retention, though, the root user cannot log in to the appliance except via the system console. Therefore, it is important to assign this authorization to the appropriate user role.
- `retentionPeriods` - User can modify the retention periods, except the grace period, for both mandatory and privileged file retention. Although a file cannot be retained for less than the minimum period, a user could change the other periods such that the file can be deleted earlier than originally set.
- `retentionAuto` - For both mandatory and privileged file retention, user can modify the grace retention period, which controls automatic file retention. When the grace period expires and the file has not been modified during the grace period, the file is automatically retained at the default retention period setting.
- `retentionOnexpiry` - For both mandatory and privileged file retention, user can set the behavior for files after their retention has expired by setting property `retention.policy.onexpiry`. The property can be set to `off` (default), `delete`, or `hold`.

  This authorization and its accompanying property are supported in deferred update File Retention on Expiry, which is available with software release OS8.8.63 or later. For more information, see File Retention on Expiry Policy.

Note that no special authorization is required to create a share with privileged file retention.

Configure the System for Mandatory File Retention

Before mandatory file retention can be used, the appliance must be configured properly, and the System Settings' Retention service must be enabled. Configure the appliance and enable the service by following these steps:

1. Ensure that the NTP service is enabled and that its property `Sync all time offsets` is enabled. While the system contains shares with mandatory retention, the NTP service must remain enabled with at least one responding NTP server. All changes to the list of servers are validated.
2. In the SSH service, set property `permit_root_login` to `false`. This disallows the root user to log in via the CLI, except for system console operations. This property cannot be modified while the system contains shares with mandatory retention enabled.
3. In the HTTPS service, set property `Permit Root Login to BUI/REST` to `false`. This disallows the root user to log in via the BUI or RESTful API. This property cannot be modified while the system contains a mandatorily retained file.
4. Enable the **System Settings' Retention** service.

If the system is not properly configured and the retention service enabled, no new mandatory file retention can be created.

For information on appliance service viewing, enabling/disabling, and setting properties, see Managing Services.

Configuring NFS to Allow Retained File Deletion (Optional)

Optionally, you can configure the NFS service to give root users (UID 0) on other systems the ability to delete files before their retention has expired.

The filesystem should be shared with root access to an appropriate, and ideally very limited, set of client systems. The root user on those client systems will be able to delete retained files when the filesystem retention policy is privileged, but not mandatory. For information on setting up these NFS client systems and granting root access, see NFS Protocol Share Mode Exceptions.

For non-Solaris NFS clients, the retention time can be displayed in the `atime` field by disabling `atime` updates on the share: `set atime=false`. Solaris clients can see the `retentiontime` by running `ls -l%all <filename>`.

Configuring SMB to Allow Retained File Deletion (Optional)

SMB users in the Administrators group can delete retained files before the expiry date when property `retention.policy` is set to `privileged`. Users can be added either to the Administrators group in Active Directory or to a local SMB group as an Administrator. For information on local SMB groups, see Adding a User to an SMB Local Group.

Planning Guidelines for File Retention

When planning for file retention, observe the following guidelines:

- Mandatory file retention affects the filesystem, project, and storage pool. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled. Note that to rename a storage pool, you must unconfigure it and then immediately import it with a new name. You cannot unconfigure a storage pool with mandatorily retained datasets.
- If you think there is a risk of someone unconfiguring a storage pool to destroy it at the same time that someone else is wrongly adding a mandatorily retained filesystem or project to the storage pool, set the maximum file retention period to a higher value.
- For mandatory file retention, the storage pool profile must provide redundancy. Therefore, the striped profile cannot be used with storage pools with mandatorily retained files.
- Directories cannot be renamed until all retained files within them have been deleted or moved to another directory. This preserves the name of the retained file, including its path.

File retention affects filesystem functionality in the following areas:

- **Editing** - Although a retained file cannot be modified, a user with role authorization `retentionPeriods` can edit a filesystem to change the retention periods only, not including the grace period. To edit the grace period, a user must have role authorization `retentionAuto`.
- **Deleting** - A filesystem with privileged retention can be deleted at any time, even if unexpired files exist. A filesystem with mandatory retention can be deleted after all of its files have expired. A file with the file retention on expiry policy set to `hold` can be deleted when the policy is set to other than `hold`.
- **Moving** - A filesystem with privileged or mandatory file retention can be moved to another project.

- **Renaming** - A file or project with privileged or mandatory file retention cannot be renamed, neither when file retention is active nor after the retention has expired.
- **Snapshots:**
  - **Editing, Deleting, Moving, Renaming** - The same principles apply as for retained files and projects not within a snapshot.
  - **Cloning** - A snapshot containing a filesystem with retention can be cloned. A file with mandatory retention protects its ancestors and clone descendants from destruction. The `hold` setting for the file retention on expiry policy applies to clones.
  - **Rollback** - Rollback can be performed on a filesystem with the privileged retention policy, even when unexpired retained files exist. Filesystems with the mandatory retention policy can never be rolled back, even when all retained files have expired.
- **Cloud Backup** - When a retained filesystem or project is used as the snapshot for a cloud backup, the same retention principles apply. Additionally, snapshots, themselves, can support snapshot retention. For more information, see Taking a Snapshot - BUI, CLI.
- **Remote Replication** - Snapshots of a retained filesystem or project have the same constraints as the file retention feature. Also, the parent and children of a snapshot containing a retained filesystem or project follow the same rules.

  When replicating to a target appliance, that appliance must support the file retention feature by accepting the File Retention deferred update, available with software version OS8.8.45 or later. When replicating to a different storage pool on the same appliance, the target pool must have a redundant profile for a file with mandatory retention. Therefore, the striped profile cannot be used with storage pools with mandatorily retained files. When replicating to an NFS server for offline replication, file retention is maintained.

  Reverse replication is not supported to a filesystem with mandatory retention. This action would require rolling back the filesystem prior to the replication reversal, which is not allowed.

  The `hold` setting for the file retention on expiry policy does not affect filesystems within a replication package nor does it affect the storage pool containing such replication packages.
- **Appliance Factory Reset** - The appliance factory reset feature is disabled if actively retained filesystems exist with mandatory file retention.

Creating a Filesystem or Project with File Retention

File retention is set at file creation, and a filesystem inherits this setting from its project if retention was set at the project level. Therefore, set retention when creating a new project or when creating a new filesystem within an unretained project. The available settings for the file retention policy are as follows:

- **Disabled** - No file retention policy is set. This is the default setting.
- **Privileged override** - Allows a user with the `retentionPeriods` role authorization to override the retention periods, except the grace period, for both mandatory and privileged file retention. Although a file cannot be retained for less than the minimum period, a user could change the other periods backwards, such that the file could be deleted earlier than originally set. The retention period can also be extended, but not beyond the maximum retention period setting.
- **Mandatory (no override)** - A filesystem with mandatory retention cannot be deleted before the retention period expires, and the retention period cannot be modified or overridden.

While creating file retention, set the retention period properties described in Retention Period Settings. In the BUI, these properties are located under the **ACCESS** tab. In the CLI, use the `get` command at the filesystem or project level to see the retention period properties. Set a value and a time measurement; no entry is the same as a zero-value entry.

Optionally, set the file retention on expiry policy at the project or filesystem level. This policy allows automatic deletion after file retention expiry or sets the file retention to an indefinite hold after expiry. In addition, the property `retention.period.deletegrace` can be set to the amount of time to delay automatically deleting the files. To remove a hold, set the policy to other than `hold`. The default setting is `off` and no action occurs after the file's retention expires. User authorization `retentionOnexpiry` is required to set property `retention.policy.onexpiry`. In the BUI, this property is located under the **ACCESS** tab. In the CLI, use the `get` command at the filesystem or project level to see the policy property. For information, see File Retention on Expiry Policy.

Optionally, set the file ACL/permission changes policy at the project or filesystem level. In the BUI, the property is located under the **ACCESS** tab. In the CLI, use the `get` command at the filesystem or project level to see the policy property. When set to `on`, property `retention.policy.changeacl` allows you to change ACL settings or file permissions, other than write, on a retained file. Thus, if the ACL or permissions were set incorrectly when a file was initially retained, this policy allows you to change those settings.

Manual file retention is accomplished in one of four ways:

- The file is retained after the minimum retention period if no maximum and/or default period was set.
- The file is retained after the maximum retention period if no minimum and/or default period was set.
- The file is retained after the default retention period if a minimum and/or maximum period was not set.
- For a file with automatic file retention set, the file can be manually retained if the minimum retention period is set to a lower value than the (automatic file retention) grace period, and the default period is set to a higher value than the grace period.

Viewing the Retention Policy Type and Statistics

To view the retention policy type and statistics, use the following procedures. Note that in the BUI and if a dataset cannot be destroyed, its trash icon is not highlighted.

- **Filesystem** - In the shares area of the software, navigate to the filesystem's project and then the filesystem. In the BUI, additionally select the **ACCESS** tab. In the CLI, issue the `get` command. The following file retention attributes are displayed:
  - Policy type
  - Minimum period
  - Maximum period
  - Default period
  - Grace retention period
  - Grace delete period
  - Number of filesystems retained
  - Retention status: expiration date, time, and if expired
  - On expiry policy

- **Project** - In the shares area of the software, navigate to the project. In the CLI, issue the `get` command. In the BUI, the policy type is displayed in the Static Properties column. Select the **ACCESS** tab to view the retention period. The following file retention expiration attributes are displayed for each filesystem:
  - Date
  - Time
  - If expired
  - Policy type: In the BUI, a solid locked icon indicates mandatory file retention, and a non-solid unlocked icon indicates privileged file retention.
- **Storage Pool** - In the configuration storage area of the software, navigate to the storage pool. In the CLI, issue the `get` command. The following file retention attributes are displayed:
  - Policy type is always mandatory; privileged file retention does not affect a storage pool
  - Retention status: expiration date, time, and if expired
  - Number of filesystems retained
  - Number of filesystems set to `hold`

Related Topics

- Creating a Filesystem - BUI, CLI
- Creating a Project - BUI, CLI
- Configuring a Storage Pool - BUI, CLI
- Project and Share Properties
- Editing Authorizations for a Role - BUI, CLI
- Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*

Project and Share Properties

All projects and shares have a number of associated properties which can be set using the BUI or CLI. For a list of property names and descriptions, select one of these links:

- Inherited Properties
- Project Properties
- Filesystem Properties
- LUN Properties

Project and share properties can be one of the following types:

Table 4-1 Project and Share Property Types

Property Type	Description
Inherited	Inherited properties, the most common property type, represent most of the configurable project and share properties. Shares that are part of a project can either have local settings for properties, or they can inherit their settings from the parent project. By default, shares inherit all properties from the project. If a property is changed on a project, all shares that inherit that property are updated to reflect the new value. When inherited, all properties have the same value as the parent project, with the exception of the mountpoint and SMB properties. When inherited, these properties concatenate the project setting with their own share name.
Read only	Read-only properties represent statistics about the project and share and cannot be changed. The most common properties of this type are space usage statistics.
Space management	Space management properties (quota and reservation) apply to both shares and projects, but are not inherited. A project with a quota of 100G will be enforced across all shares, but each individual share will have no quota unless explicitly set.
Static (Create time)	Static properties are specified at filesystem or LUN creation time, but cannot be changed once the share has been created. These properties control the on-disk data structures, and include internationalization settings, case sensitivity, and volume block size.
Project default	Project default properties are set on a project, but do not affect the project itself. They are used to populate the initial settings when creating a filesystem or LUN, and can be useful when shares have a common set of non-inheritable properties. Changing these properties do not affect existing shares, and the properties can be changed before or after creating the share.
Filesystem local	Filesystem local properties apply only to filesystems, and are convenience properties for managing the root directory of the filesystem. They cannot be set on projects. These access control properties can also be set by in-band protocol operations.
LUN local	LUN local properties apply only to LUNs and are not inherited. They cannot be set on projects.
Custom	Custom properties are user-defined properties.

Inherited Properties

Inherited properties are properties that inherit their value from their parent. Projects can inherit properties from their parent pool, and shares can inherit properties from their parent project. Usually, properties that can be inherited can alternatively be set explicitly on a project or share. When a value is explicitly set for a property that was initially inherited, that property is then a local property.

The BUI only allows the properties to be inherited all at once, while the CLI allows for individual properties to be inherited.

By default, shares inherit all of their properties from their parent project. If a property is changed on a project, all shares that inherit that property (that property has not been set locally) are updated to reflect the new value. When inherited, all properties have the same value as the parent project, with the exception of the mountpoint and SMB properties. When inherited, the mountpoint and SMB properties concatenate the project setting with their own share name.

In addition to the properties described in the following sections, encryption properties also are inherited. If a pool is encrypted, all projects created under that pool inherit the pool's encryption settings. If a project is encrypted, all shares created under that project inherit the project's encryption settings. See Encryption Properties. Creating an Encrypted Project - BUI, CLI, and Creating an Encrypted Filesystem or LUN - BUI, CLI.

For a description of how inherited property settings could be changed by restoration from backup, and how to avoid that, see Inherited Property Settings after Restoration from Backup.

Mountpoint

The mountpoint property is the location where the filesystem is mounted. This property is only valid for filesystems.

The following restrictions apply to the mountpoint property:

- Must be under /export
- Cannot conflict with another share
- Cannot conflict with another share on cluster peer to allow for proper failover

When inheriting the mountpoint property, the current dataset name is appended to the project's mountpoint setting, joined with a forward slash (/). For example, if the `home` project has the mountpoint setting `/export/home`, then `home/chris` would inherit the mountpoint `/export/home/chris`.

SMB shares are exported via their resource name, and the mountpoint is not visible over the protocol. However, even SMB-only shares must have a valid unique mountpoint on the appliance.

Mountpoints can be nested underneath other shares' mountpoints, although this has some limitations. For more information, see Working with Filesystem Namespace.

Read Only

The `read-only` property controls whether the filesystem contents are read only. This property is only valid for filesystems.

The contents of a read-only filesystem cannot be modified, regardless of any protocol settings. This setting does not affect the ability to rename, destroy, or change properties of the filesystem. In addition, when a filesystem is read only, access control properties cannot be altered, because they require modifying the attributes of the root directory of the filesystem.

Update Access Time on Read

The update access time on read property controls whether the access time for files is updated on read. This property is only valid for filesystems.

POSIX standards require that the access time for a file properly reflect the last time it was read. This requires issuing writes to the underlying filesystem even for a mostly read-only workload. For working sets consisting primarily of reads over a large number of files, turning off this property may yield performance improvements at the expense of standards conformance. These updates happen asynchronously and are grouped together, so its effect should not be visible except under heavy load.

Non-Blocking Mandatory Locking

The `non-blocking mandatory locking` property controls whether SMB locking semantics are enforced over POSIX semantics. This property is only valid for filesystems.

By default, filesystems implement file behavior according to POSIX standards. These standards are fundamentally incompatible with the behavior required by the SMB protocol. For shares where the primary protocol is SMB, this option should always be enabled. Changing this property requires all clients to be disconnected and reconnect.

Data Deduplication

The data deduplication property controls whether duplicate copies of data are eliminated. Deduplication is synchronous, pool-wide, block-based, and can be enabled on a per project or share basis.

Before deduplication can be enabled on a project or share, configure the storage pool with meta devices. Meta devices are designated cache devices used to store specific types of metadata to optimize use cases like deduplication.

Deduplication is also only available on datasets with a record size 128K or above.

To enable deduplication, select the Data Deduplication checkbox on the general properties screen for projects or shares. The size of the deduplicated data, as well as the deduplication ratio, will appear in the usage area of the Status Dashboard. Data written with deduplication enabled is entered into the deduplication table indexed by the data checksum. Deduplication forces the use of the cryptographically strong SHA-256 checksum. Subsequent writes will identify duplicate data and retain only the existing copy on disk. Deduplication can only happen between blocks of the same size, data written with the same record size. For best results, set the record size to that of the application using the data; for streaming workloads, use a large record size.

> **Note:**
>
> Starting with the Data Deduplication deferred update in OS8.7, if replication is configured for a deduplicated project or share, a compatibility test will run to determine if the replication target has the required software and meta device for receiving deduplicated updates. If the target is running OS8.6 or earlier, replication updates with deduplication enabled will fail, and an alert will post indicating that data deduplication needs to be disabled at the source. If the target is running OS8.7, but does not have the required meta device, the target will ignore the incoming data deduplication property, and an alert will post to show that the target system will deliberately ignore the deduplication settings during replication receive. If a OS8.6 or earlier source has been replicating to a target with deduplication enabled, the compatibility test will check the target for a meta device or deduplicated share in the package. If the compatibility test finds either of these, the target will preserve the deduplication settings when receiving replication updates.

If your data does not contain any duplicates, enabling data deduplication will add overhead (a more CPU-intensive checksum and on-disk deduplication table entries) without providing any benefit. If your data does contain duplicates, enabling data deduplication will both save space by storing only one copy of a given block regardless of how many times it occurs. Deduplication could impact performance in that the checksum is more expensive to compute and the metadata of the deduplication table must be accessed and maintained.

Note that deduplication has no effect on the calculated size of a share and does not affect the amount of space used for the pool. For example, if two shares contain the same 512 GB file, each will appear to be 512 GB in size, but the total for the pool will be just 512 GB, and deduplication will be reported as 512G (2x). If three shares contain the same 512 GB file, each appears as 512 GB in size, the total for the pool will be 512 GB, and deduplication will be 1024G (3x).

There are 3 sets of analytics used to monitor performance of deduplication:

- **ZFS DMU operations (by DMU object type)** - This analytic will show you how many operations are being performed against the Data Deduplication Table compared to other ZFS operations.
- **Meta device bytes used (by pool)** - Amount of space used on the metadata devices.

  This statistic will remain blank until at least 1% of the meta device capacity is used.
- **Meta device percent used (by pool)** - Percent of space used on the metadata devices.

  This statistic will remain blank until at least 1% of the meta device capacity is used.

To use deduplication with encryption, keep in mind that only AES with the CCM mode encryption is compatible with deduplication. For more information, see Managing Encryption Keys.

Data Compression

The data compression property controls whether data is compressed before being written to disk. Shares can optionally compress data before writing to the storage pool. This allows for much greater storage utilization at the expense of increased CPU utilization. By default, no compression is done. If the compression does not yield a minimum space savings, it is not committed to disk to avoid unnecessary decompression when reading back the data. Before choosing a compression algorithm, it is recommended that you perform any necessary performance tests and measure the achieved compression ratio.

BUI value	CLI value	Description
Off	`off`	No compression is done.
LZ4	`lz4`	An algorithm that typically consumes less CPU than GZIP-2, but compresses better than LZJB, depending on the data that is compressed.
LZJB (Fastest)	`lzjb`	A simple run-length encoding that only works for sufficiently simple inputs, but does not consume much CPU.
GZIP-2 (Fast)	`gzip-2`	A lightweight version of the gzip compression algorithm.
GZIP (Default)	`gzip`	The standard gzip compression algorithm.
GZIP-9 (Best Compression)	`gzip-9`	Highest achievable compression using gzip. This consumes a significant amount of CPU and can often yield only marginal gains.

Checksum

The checksum property controls the checksum used for data blocks. On the appliance, all data is checksummed on disk, and in such a way to avoid traditional pitfalls (phantom reads and write in particular). This allows the system to detect invalid data returned from the devices. The default checksum (Fletcher4) is sufficient for normal operation, but users can increase the checksum strength at the expense of additional CPU load. Metadata is always checksummed using the same algorithm, so this only affects user data (files or LUN blocks).

BUI value	CLI value	Description
Fletcher 2 (Legacy)	`fletcher2`	16-bit fletcher checksum
Fletcher 4 (Standard)	`fletcher4`	32-bit fletcher checksum
SHA-256 (Extra Strong)	`sha256`	SHA-256 checksum
SHA-256-MAC	`sha256mac`	

Cache Device Usage

The cache device usage property controls whether cache devices are used for the share. By default, all datasets make use of any cache devices on the system. Cache devices are configured as part of the storage pool and provide an extra layer of caching for faster tiered access. For more information on cache devices, see Configuring Storage. This property is independent of whether there are any cache devices currently configured in the storage pool. For example, it is possible to have this property set to `all` even if there are no cache devices present. If any such devices are added in the future, the share will automatically take advantage of the additional performance. This property does not affect use of the primary (DRAM) cache.

BUI Value	CLI Value	Description
All data and metadata	`all`	All normal file or LUN data is cached, as well as any metadata.
Metadata only	`metadata`	Only metadata is kept on cache devices. This allows for rapid traversal of directory structures, but retrieving file contents may require reading from the data devices.
Do not use cache devices	`none`	No data in this share is cached on the cache device. Data is only cached in the primary cache or stored on data devices.

Synchronous Write Bias

The synchronous write bias property controls the behavior when servicing synchronous writes. By default, the system optimizes synchronous writes for latency, which leverages the log devices to provide fast response times. In a system with multiple disjointed filesystems, this can cause contention on the log devices that can increase latency across all consumers. Even with multiple filesystems requesting synchronous semantics, it may be the case that some filesystems are more latency-sensitive than others.

A common case is a database that has a separate log. The log is extremely latency sensitive, and while the database itself also requires synchronous semantics, it is heavier bandwidth and not latency sensitive. In this environment, setting this property to `throughput` on the main database while leaving the log filesystem as `latency` can result in significant performance improvements. This setting will change behavior even when no log devices are present, though the effects may be less dramatic.

The synchronous write bias setting can be bypassed by the Oracle Intelligent Storage Protocol. Instead of using the write bias defined in the file system, Oracle Intelligent Storage Protocol can use the write bias value provided by the Oracle Database NFSv4.0 or NFSv4.1 client. The write bias value sent by the Oracle Database NFSv4.0 or NFSv4.1 client is used only for that write request.

BUI Value	CLI Value	Description
Latency	`latency`	Synchronous writes are optimized for latency, leveraging the dedicated log device(s), if any.
Throughput	`throughput`	Synchronous writes are optimized for throughput. Data is written to the primary data disks instead of the log device(s), and the writes are performed in a way that optimizes for total bandwidth of the system. Log devices will be used for small amounts of metadata associated with the data writes.

Database Record Size

The database record size property specifies a suggested block size for files in a filesystem. This property is only valid for filesystems and is designed for use with database workloads that access files in fixed-size records. The system automatically tunes block sizes according to internal algorithms optimized for typical access patterns. For databases that create very large files but access them in small random chunks, these algorithms may be suboptimal. Specifying a record size greater than or equal to the record size of the database can result in significant performance gains. Use of this property for general purpose filesystems is strongly discouraged, and may adversely affect performance.

The default record size is 128 KB. The size specified must be a power of two greater than or equal to 512 and less than or equal to 1 MB. Changing the filesystem's record size affects only files created afterward; existing files and received data are unaffected. If block sizes greater than 128K are used for projects or shares, replication of those projects or shares to systems that do not support large block sizes will fail.

The database record size setting can be bypassed by Oracle Intelligent Storage Protocol. Instead of using the record size defined in the filesystem, Oracle Intelligent Storage Protocol can use the block-size value provided by the Oracle Database NFSv4.0 or NFSv4.1 client. The block size provided by the Oracle Database NFSv4.0 or NFSv4.1 client can only be applied when creating a new database files or table. Block sizes of existing files and tables will not be changed. For more information, see Configuring Oracle ZFS Storage Appliance for Oracle Database Clients.

Additional Replication

The additional replication property controls the number of copies stored of each block, above and beyond any redundancy of the storage pool. Metadata is always stored with multiple copies, but this property allows the same behavior to be applied to data blocks. The storage pool attempts to store these extra blocks on different devices, but it is not guaranteed. In addition, a storage pool cannot be imported if a complete logical device (RAID stripe, mirrored pair, and so on) is lost. This property is not a replacement for proper replication in the storage pool, but can be reassuring for administrators.

Virus Scan

The virus scan property controls whether a filesystem is scanned for viruses. This property is only valid for filesystems. This property setting is independent of the state of the virus scan service. Even if the Virus Scan service is enabled, filesystem scanning must be explicitly enabled using this property. Similarly, virus scanning can be enabled for a particular share even if the service itself is off. For more information about configuration virus scanning, see Virus Scan Configuration.

Prevent Destruction

When set, the share or project cannot be destroyed. This includes destroying a share through dependent clones, destroying a share within a project, or destroying a replication package. However, it does not affect shares destroyed through replication updates. If a share is destroyed on an Oracle ZFS Storage Appliance system that is the source for replication, the corresponding share on the target will be destroyed, even if this property is set. To destroy the share, the property must first be explicitly turned off as a separate step. This property is off by default.

Restrict Ownership Change

By default, ownership of files cannot be changed except by a root user (on a suitable client with a root-enabled export). This property can be turned off on a per-filesystem or per-project basis by turning off this property. When off, file ownership can be changed by the owner of the file or directory, effectively allowing users to "give away" their own files. When ownership is changed, any `setuid` or `setgid` bits are stripped, preventing users from escalating privileges through this operation.

LUN Local Properties

These properties apply only to LUNs and are not inherited. They cannot be set on projects.

Volume Size

The volume size property is the logical size of the LUN as exported over iSCSI. This property controls the size of the LUN. By default, LUNs reserve enough space to completely fill the volume. Changing the size of a LUN while actively exported to clients may yield undefined results. It may require clients to reconnect and/or cause data corruption on the filesystem on top of the LUN. Check best practices for your particular iSCSI client before attempting this operation.

Thin Provisioned

The thin provisioned property controls whether space is reserved for the volume. By default, a LUN reserves exactly enough space to completely fill the volume. This ensures that clients will not get out-of-space errors at inopportune times. This property allows the volume size to exceed the amount of available space. When set, the LUN will consume only the space that has been written to the LUN. While this allows for thin provisioning of LUNs, most filesystems do not expect to get "out of space" from underlying devices, and if the share runs out of space, it may cause instability and/or data corruption on clients.

When not set, the volume size behaves like a reservation excluding snapshots. It therefore has the same pathologies, including failure to take snapshots if the snapshot could theoretically diverge to the point of exceeding the amount of available space. For more information, see the `Reservation` property in Managing Filesystem and Project Space.

Thin provisioned LUNs can optionally take advantage of the space reclamation feature that returns free space to the storage pool, as described in Space Reclamation.

The logical block provisioning (LBP) threshold for thin provisioned LUNs is another way to manage space. The `LBPthreshold` command is available in the RESTful API only, and it sets the LBP threshold for thin provisioned LUNs within a specified storage pool. When set to the default value of 0, the LBP threshold is disabled. To set the threshold, specify the average storage consumption rate. When the threshold value is exceeded, an error message is returned, and you can manage the space accordingly. See also Set LBP Threshold in *Oracle ZFS Storage Appliance RESTful API Guide, Release OS8.8.x*.

Space Reclamation

The `space reclamation` property controls whether LUN free space is returned to the storage pool. When set to true, the client operating system issues the `SCSI UNMAP` command to return unused logically provisioned space after volume data blocks have been deleted by the host operating system. Therefore, this property must be used in conjunction with the `thin provisioned` property, which logically sets the volume or LUN size, and it allows the size to exceed the amount of available space, which can be considered free space.

Volume Block Size

The volume block size property sets the native block size for LUNs. This can be any power of 2 from 512 bytes to 1 M, and the default is 8 K. This property is static; it is set when the LUN is created and cannot be changed.

> **Note:**
>
> LUNs with a volume block size smaller than 4 K may cause performance degradation.

Other Properties

The following "other" properties are available: Project Default, Filesystem Local, Space Management, Read-only, and Custom.

Project Default

The project default properties are set on a project, but do not affect the project itself. They are used to populate the initial settings when creating a filesystem or LUN, and can be useful when shares have a common set of non-inheritable properties. Changing these properties do not affect existing shares, and the properties can be changed before or after creating the share.

Filesystem Local

The filesystem local properties apply only to filesystems, and are convenience properties for managing the root directory of the filesystem. They are not inherited and cannot be set on projects. These access control properties can also be set by in-band protocol operations.

Space Management

The space management properties (quota and reservation) apply to both shares and projects, but are not inherited. A project with a quota of 100 G will be enforced across all shares, but each individual share will have no quota unless explicitly set.

Read Only

The read only properties represent statistics about the project and share and cannot be changed. The most common properties of this type are space usage statistics.

Custom

Custom properties are user-defined using a schema. For more information, see Working with Schemas.

Static Properties

Static (create-time) properties are specified at filesystem or LUN creation time, but cannot be changed after the share has been created. These properties control the on-disk data structures, and include internationalization settings, case sensitivity, and volume block size.

In the BUI, static properties can be viewed on the left side of the interface when editing a filesystem or LUN.

Table 4-2 Filesystem and LUN Static Properties

BUI Name	CLI Name	Description
Creation date	`creation`	Indicates the date of creation.
Compression ratio	`compressratio`	Current compression ratio for the filesystem or LUN, which is a product of the compression algorithm. For more information, see Compression ratio.
Case sensitivity	`casesensitivity`	The case sensitivity property controls whether directory lookups are case-sensitive or case-insensitive. For more information, see Case sensitivity.
Reject non UTF-8	`utf8only`	This property enforces UTF-8 encoding for all files and directories. For more information, see Reject non UTF-8.
Normalization	`normalization`	The normalization property controls what unicode normalization, if any, is performed on filesystems and directories. Unicode supports the ability to have the same logical name represented by different encodings. For more information, see Normalization.
Volume block size (LUNs only)	`volblocksize`	The volume block size property sets the native block size for LUNs. For more information, see Volume block size.
Origin	`origin`	Shows the name of the snapshot from which it was cloned. For more information, see Origin.
Data migration source (Filesystems only)	`shadow`	Location of the source if the filesystem is actively shadowing an existing filesystem, either locally or over NFS. For more information, see Data Migration Source

Compression Ratio

If compression is enabled, this property shows the compression ratio currently achieved for the share. This is expressed as a multiplier. For example, a compression of 2x means that the data is consuming half as much space as the uncompressed contents. For more information about selecting a compression algorithm, see "Data Compression" described in Inherited Properties.

Case Sensitivity

The case sensitivity property controls whether directory lookups are case-sensitive or case-insensitive. It supports the following options:

BUI Value	CLI Value	Description
Mixed	`mixed`	Case sensitivity depends on the protocol being used. For NFS, FTP, and HTTP, lookups are case-sensitive. For SMB, lookups are case-insensitive. This is default, and prioritizes conformance of the various protocols over cross-protocol consistency. When using this mode, it is possible to create files that are distinct over case-sensitive protocols, but clash when accessed over SMB. In this situation, the SMB server will create a "mangled" version of the conflicts that uniquely identify the filename.
Insensitive	`insensitive`	All lookups are case-insensitive, even over protocols (such as NFS) that are traditionally case-sensitive. This can cause confusion for clients of these protocols, but prevents clients from creating name conflicts that would cause mangled names to be used over SMB. This setting should only be used where SMB is the primary protocol and alternative protocols are considered second-class, where conformance to expected standards is not an issue.
Sensitive	`sensitive`	All lookups are case-sensitive, even over SMB where lookups are traditionally case-insensitive. In general, this setting should not be used because the SMB server can deal with name conflicts via mangled names, and may cause Windows applications to behave strangely.

Reject non UTF-8

This property enforces UTF-8 encoding for all files and directories. When set, attempts to create a file or directory with an invalid UTF-8 encoding will fail. This only affects NFSv3, where the encoding is not defined by the standard. NFSv4.0 and NFSv4.1 always use UTF-8, and SMB negotiates the appropriate encoding. This setting should normally be `on`, or else SMB (which must know the encoding for case sensitive comparisons, among other operations) will be unable to decode filenames that are created with and invalid UTF-8 encoding. This setting should only be set to `off` in pre-existing NFSv3 deployments where clients are configured to use different encodings. Enabling SMB, NFSv4.0 or NFSv4.1 when this property is set to `off` can yield undefined results if a NFSv3 client creates a file or directory that is not a valid UTF-8 encoding. This property must be set to `on` if the normalization property is set to anything other than `none`.

Normalization

The normalization property controls which unicode normalization, if any, is performed on filesystems and directories. Unicode supports the ability to have the same logical name represented by different encodings. Without normalization, the on-disk name stored will be different, and lookups using one of the alternative forms will fail depending on how the file was created and how it is accessed. If this property is set to anything other than `none` (the default), the `Reject non UTF-8` property must also be set to `on`.

BUI Value	CLI Value	Description
None	`none`	No normalization is done.
Form C	`formC`	**Normalization Form Canonical Composition (NFC)** - Characters are decomposed and then recomposed by canonical equivalence.

BUI Value	CLI Value	Description
Form D	`formD`	**Normalization Form Canonical Decomposition (NFD)** - Characters are decomposed by canonical equivalence.
Form KC	`formKC`	**Normalization Form Compatibility Composition (NFKC)** - Characters are decomposed by compatibility equivalence, then recomposed by canonical equivalence.
Form KD	`formKD`	**Normalization Form Compatibility Decomposition (NFKD)** - Characters are decomposed by compatibility equivalence.

Volume Block Size

The volume block size property sets the native block size for LUNs. This can be any power of 2 from 512 bytes to 1 M, and the default is 8 K.

> **Note:**
>
> LUNs with a volume block size smaller than 4 K may cause performance degradation.

Origin

If this is a clone, this is the name of the snapshot from which it was cloned.

Data Migration Source

If set, then this filesystem is actively shadowing an existing filesystem, either locally or over NFS. For more information about data migration, see Shadow Migration.

Project Properties

In the CLI, use the `get` command to see a list of all properties. Use the `list` command to list all children.

The following table shows **Create Project** properties.

Table 4-3 Create Project Properties

BUI Name	CLI Name	Property Type	Description
Name	`project`	Static	Defines the name of the project.
Encryption	`encryption`	Inherited	Defines the encryption type. For more information see, Managing Encryption Keys.
Inherit key	`--`	--	If selected, indicates that the encryption key is inherited from the parent pool.
Key	`key`	Inherited	Sets a specific encryption key that is used when the key is not inherited from the parent pool.
Keyname	`keyname`	Static	Identifies the key.

Table 4-3   (Cont.) Create Project Properties

BUI Name	CLI Name	Property Type	Description
File retention policy	`retention.policy`	Inherited	Sets one of the following options: • Disabled (default): No file retention policy is set. • Privileged override: Sets the privileged file retention policy. • Mandatory (no override): Sets the mandatory file retention policy. Project filesystems inherit file retention properties when the project is created, and the properties are local to the project's filesystems. Therefore, if a project is later renamed, the filesystems still retain their originally inherited file retention properties.
File retention on expiry policy	`retention.policy.onexpiry`	Inherited	Determines behavior when file retention expires. Sets one of the following options: • Off (default): The file is not affected, and it remains on the system after retention expires. • Delete: The file is deleted after expiration has been met. • Hold: The file transitions to an indefinite hold. The file cannot be deleted until set to `off` or `delete`.
Delete expired files after	`retention.period.deletegrace`	Inherited	Number of seconds/hours/days/years that automatic file deletion is delayed when the file retention on expiry policy is set to `delete`. Default value: 0 days.
Allow permission changes on retained files	`retention.policy.changeacl`	Inherited	Determines if a retained file's ACL/ permissions can be changed. Default value is `off` and the file's ACL/ permissions cannot be changed. When set to `on`, the ACL settings/permissions, other than write, can be changed on a retained file.

The following table shows **General - Space Usage - Data** project properties.

Table 4-4   General - Space Usage - Data Project Properties

BUI Name	CLI Name	Property Type	Description
Quota	`quota`	Space management	Sets a limit on the amount of space that can be consumed by any particular entity.
Reservation	`reservation`	Space management	Represents a guarantee of space that can be consumed by any particular entity.

The following table shows **General - Space Usage - Users & Groups** project properties.

Table 4-5 General - Space Usage - Users & Groups Project Properties

BUI Name	CLI Name	Property Type	Description
Default user quota	defaultuserquota	Space management	Sets a limit on the amount of space that can be consumed by the user.
Default group quota	defaultgroupquota	Space management	Sets a limit on the amount of space that can be consumed by the group.
User and group	users / groups	--	Specifies users and/or groups.
Usage	--	Space management	Shows the amount of data used by users and/or groups.

The following table shows **General - Inherited Properties** project properties.

Table 4-6 General - Inherited Project Properties

BUI Name	CLI Name	Property Type	Description
Mountpoint	mountpoint	Inherited	Controls the path used to export filesystems. For more information, see Mountpoint.
Read only	readonly	Inherited	Controls whether the filesystem contents are read only. For more information, see Read only.
Update access time on read	atime	Inherited	Controls whether the access time for files is updated on read. For more information, see Update access time on read.
Non-blocking mandatory locking	nbmand	inherited	Controls whether SMB locking semantics are enforced over POSIX semantics. For more information, see Non-blocking mandatory locking.
Data deduplication	dedup	Inherited	Controls whether duplicate copies of data are eliminated. For more information, see Data Deduplication.
Data compression	compression	Inherited	Controls whether data is compressed before being written to disk. For more information, see Data Compression.
Checksum	checksum	Inherited	Controls the checksum used for data blocks. For more information, see Checksum.
Cache device usage	secondarycache	Inherited	Controls whether cache devices are used for the share. For more information, see Cache device usage.
Synchronous write bias	logbias	Inherited	Controls the behavior when servicing synchronous writes. For more information, see Synchronous write bias.
Database record size	recordsize	Inherited	Specifies a suggested block size for files in the filesystem. For more information, see Database record size.

Table 4-6 (Cont.) General - Inherited Project Properties

BUI Name	CLI Name	Property Type	Description
Additional Replication	`copies`	Inherited	Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. For more information, see Additional replication.
Virus scan	`vscan`	Inherited	Controls whether a filesystem is scanned for viruses. For more information, see Virus scan.
Prevent destruction	`nodestroy`	Inherited	Prevents shares or projects from being destroyed when set. For more information, see Prevent destruction. For preventing destruction at the storage pool level, see Destroy Prevention and Approval.
Restrict ownership change	`rstchown`	Inherited	Controls the ownership and can be turned off on a per-filesystem or per-project basis. For more information, see Restrict ownership change.

The following table shows **General - Custom Properties** project properties.

Table 4-7 General - Custom Project Properties

BUI Name	CLI Name	Property Type	Description
Schema	`custom`	--	Custom properties can be added as needed to attach user-defined tags to projects and shares. For more information, see Schema Properties.

The following table shows **General - Default Settings - Filesystems** project properties.

Table 4-8 General - Default Settings - Filesystems Project Properties

BUI Name	CLI Name	Property Type	Description
User	`default_user`	Creation default	Specifies a user ID or user name.
Group	`default_group`	Creation default	Specifies a group ID or group name.
Permissions	`default_permissions`	Creation default	Sets the default permissions for filesystem.

The following table shows **General - Default Settings - LUNs** project properties.

Table 4-9 General - Default Settings - LUNs Project Properties

BUI Name	CLI Name	Property Type	Description
Volume size	`default_volsize`	LUN only, creation default	Shows the maximum volume size and unit of measurement. For more information, see Volume size.

Table 4-9 (Cont.) General - Default Settings - LUNs Project Properties

BUI Name	CLI Name	Property Type	Description
Thin provisioned	`default_sparse`	LUN only, creation default	Indicates only the amount of space physically consumed by data is used when selected. For more information, see Thin provisioned.
Volume block size	`default_volblocksize`	Creation default	Shows the native block size for LUNs and can be set from 512 bytes to 1M; the default is 8K. For more information, see Volume block size.

The following table shows **Bandwidth** project properties.

Table 4-10 Bandwidth Project Properties

BUI Name	CLI Name	Property Type	Description
Read limit	`readlimit`	--	Sets the maximum bytes per second that can be read from a share. `M` indicates megabytes and `G` indicates gigabytes. The default setting is `unlimited`, which provides no I/O throttling.
Write limit	`writelimit`	--	Sets the maximum bytes per second that can written to a share. `M` indicates megabytes and `G` indicates gigabytes. The default setting is `unlimited`, which provides no I/O throttling.
Effective read limit	`effectivewritelimit`	--	Read-only property that reports the lowest read limit for a share.
Effective write limit	`effectivewritelimit`	--	Read-only property that reports the lowest write limit for a share.

The following table shows **Protocols - NFS** project properties.

Exceptions to the overall sharing modes may be defined for clients or collections of clients. For more information, see NFS Protocol Share Mode Exceptions.

Table 4-11 Protocols - NFS Project Properties

BUI Name	CLI Name	Property Type	Description
NFS	`sharenfs`	Inherited	NFS Protocol property settings and values are described in NFS Protocol Properties.

The following table shows **Protocols - SMB** project properties.

Exceptions to the overall sharing modes may be defined for clients or collections of clients. For more information, see SMB Protocol Share Mode Exceptions.

Table 4-12 Protocols - SMB Project Properties

BUI Name	CLI Name	Property Type	Description
SMB	`sharesmb`	Inherited	SMB Protocol property settings and values are described in SMB Protocol Properties.

The following table shows **Protocols - HTTP** (Inherit from project) project properties.

Table 4-13 Protocols - HTTP (Inherit from project) Project Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`sharedav`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Protocols - FTP** (Inherit from project) project properties.

Table 4-14 Protocols - FTP (Inherit from project) Project Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`shareftp`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Protocols - SFTP** (Inherit from project) project properties.

Table 4-15 Protocols - SFTP (Inherit from project) Project Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`sharesftp`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Protocols - TFTP** (Inherit from project) project properties.

Table 4-16 Protocols - TFTP (Inherit from project) Project Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`sharetftp`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Access** project properties.

Table 4-17 Access Project Properties

BUI Name	CLI Name	Property Type	Description
ACL behavior on mode change	`aclmode`	Inherited	Controls how a mode change request interacts with the existing ACL.
ACL inheritance behavior	`aclinherit`	Inherited	Controls how a new file or directory inherits existing ACL settings from the parent directory.

The following table shows **Snapshots - Properties** project properties.

Table 4-18 Snapshots - Properties Project Properties

BUI Name	CLI Name	Property Type	Description
.zfs/snapshot visibility	`snapdir`	Inherited	Controls whether filesystem snapshots can be accessed over data protocols at `.zfs/snapshot` in the root of the filesystem.
Scheduled snapshot label	`snaplabel`	Inherited	Appends a user-defined label to each scheduled snapshot and is blank by default.

The following table shows **Snapshots - Snapshots** project properties.

Table 4-19 Snapshots - Snapshots Project Properties

BUI Name	CLI Name	Property Type	Description
Name	`snapshot name`	--	Specifies the name of the snapshot.
Creation	`creation`	--	Specifies the date and time when the snapshot is created.
Unique	`space_unique`	--	Indicates the amount of unique space used by the snapshot.
Total	`space_data`	--	Indicates the total amount of space referenced by the snapshot.

The following table shows **Snapshots - Schedules** project properties.

Table 4-20 Snapshots - Schedules Project Properties

BUI Name	CLI Name	Property Type	Description
Frequency	`frequency`	Create time	Indicates how often the snapshot is taken.
Keep at most	`keep`	Create time	Controls the retention policy for snapshots.

The following table shows **Replication** (Inherit from project)**/Create New Actions** project properties.

Table 4-21 Replication (Inherit from project)/Create New Actions Project Properties

BUI Name	CLI Name	Property Type	Description
Target	`target`	Inherited	Identifies the replication target system.
Pool	`pool`	Inherited	Specifies the storage pool on the target where the project will be replicated.
Export data path	`export_path`	Inherited	Indicates the export data path.
Limit bandwidth	`max_bandwidth`	Inherited	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second).
Enable SSL-encryption	`use_ssl`	Inherited	Controls whether to encrypt data on the wire using SSL.
Disable compression	--	Inherited	Controls whether the compression is enabled or disabled.
Include snapshot	`include_snaps`	Inherited	Controls whether replication updates include non-replication snapshots.
Retain user snapshots on target	`retain_user_snaps_on_target`	Inherited	When set, keeps user-generated snapshots on the replication target. Continues to retain snapshots on the target until disabled.
Include clone origin as data	`include_clone_origin_as_data`	Inherited	Controls the replication of each share that was cloned from a share that is external to the replication package on the target.
Recovery point objective	`recovery_point_objective`	Inherited	Specifies the maximum tolerable amount of data loss in the event of a disaster or major outage.
Replica lag warning alert	`replica_lag_warning_alert`	Inherited	Specifies a limit, represented as a percentage of the RPO, when a minor alert is generated.
Replica lag error alert	`replica_lag_error_alert`	Inherited	Specifies a limit, represented as a percentage of the RPO, when a major alert is generated.
Update frequency	`continuous`	Inherited	Controls whether this action is being replicated continuously or at manual or scheduled intervals.

Filesystem Properties

In the CLI, use the `get` command to see a list of all properties.

The following table shows **Create Filesystem** properties.

Table 4-22 Create Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Project	`select project_name`	--	Defines which project the filesystem uses to inherit parameter settings. You can also select the default project.

Table 4-22 (Cont.) Create Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Name	`filesystem`	--	Defines the name of the filesystem.
Data migration source	`shadow`	Create time	Shows the location of the source if you are migrating data.
User	`root-user`	Filesystem local	Specifies the owner of the root directory.
Group	`root_group`	Filesystem local	Specifies the group of the root directory.
Permissions or Use Windows default permissions	`root_permission`	Filesystem local	Specifies standard UNIX permissions for the root directory, or Windows default permissions.
Inherit mountpoint	--	--	Indicates the mountpoint is inherited if selected.
Mountpoint	`mountpoint`	Inherited	Controls the path used to export filesystems. For more information, see Mountpoint.
Reject non UTF-8	`utf8only`	Create time	Enforces UTF-8 encoding for all filesystems and directories. For more information, see Reject non UTF-8.
Case sensitivity	`casesensitivity`	Create time	Controls whether directory lookups are case-sensitive, case-insensitive, or mixed. For more information, see Case sensitivity.
Normalization	`normalization`	Create time	Controls which unicode normalization, if any, is performed on filesystems and directories. For more information, see Normalization.
Encryption	`encryption`	Inherited	Defines the encryption type. For more information see, Managing Encryption Keys.
Inherit key	--	--	If selected, indicates that the encryption key is inherited from the parent project.
Key	`key`	Inherited	Sets a specific encryption key that is used when the key is not inherited from the parent project.
Keyname	`keyname`	Static	Identifies the key.
File retention policy	`retention.policy`	Inherited	Sets one of the following options: • Disabled (default): No file retention policy is set. • Privileged override: Sets the privileged file retention policy. • Mandatory (no override): Sets the mandatory file retention policy.

The following table shows **General - Space Usage - Data** filesystem properties.

Table 4-23 General - Space Usage - Data Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Quota	quota	Space management	Sets a limit on the amount of space that can be consumed by any particular entity.
Quota Include snapshots	quota_snap	Space management	Sets a limit on the amount of space that can be consumed by any particular entity including the snapshots.
Reservation	reservation	Space management	Represents a guarantee of space that can be consumed by any particular entity.
Reservation Include snapshots	reservation_snap	Space management	Represents a guarantee of space that can be consumed by any particular entity including the snapshots.

The following table shows **General - Space Usage - Users & Groups** filesystem properties.

Table 4-24 General - Space Usage - Users & Groups Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Users & Groups	--	--	Specifies the users and/or groups.
Usage	--	--	Shows the amount of data used by the users and/or groups.
Quota	quota	Space management	Sets a limit on the amount of space that can be consumed by any particular entity.

The following table shows **General - Properties** (Inherit from project) filesystem properties.

Table 4-25 General - Properties (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Mountpoint	mountpoint	Inherited	Controls the path used to export filesystems. For more information, see Mountpoint.
Read only	readonly	Inherited	Controls whether the filesystem contents are read only. For more information, see Read only.
Update access time on read	atime	Inherited	Controls whether the access time for files is updated on read. For more information, see Update access time on read.
Non-blocking mandatory locking	nbmand	inherited	Controls whether SMB locking semantics are enforced over POSIX semantics. For more information, see Non-blocking mandatory locking.
Data deduplication (warning)	dedup	Inherited	Controls whether duplicate copies of data are eliminated. For more information, see Data Deduplication.

Table 4-25 (Cont.) General - Properties (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Data compression	compression	Inherited	Controls whether data is compressed before being written to disk. For more information, see Data Compression.
Checksum	checksum	Inherited	Controls the checksum used for data blocks. For more information, see Checksum.
Cache device usage	secondarycache	Inherited	Controls whether cache devices are used for the share. For more information, see Cache device usage.
Synchronous write bias	logbias	Inherited	Controls the behavior when servicing synchronous writes. For more information, see Synchronous write bias.
Database record size	recordsize	Inherited	Specifies a suggested block size for files in the filesystem. For more information, see Database record size.
Additional Replication	copies	Inherited	Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. For more information, see Additional replication.
Virus scan	vscan	Inherited	Controls whether a filesystem is scanned for viruses. For more information, see Virus scan.
Prevent destruction	nodestroy	Inherited	Prevents shares or projects from being destroyed when set. For more information, see Prevent destruction. For preventing destruction at the storage pool level, see Destroy Prevention and Approval.
Restrict ownership change	rstchown	Inherited	Controls the ownership and can be turned off on a per-filesystem or per-project basis. For more information, see Restrict ownership change.

The following table shows **General - Custom Properties** (Inherit from Project) filesystem properties.

Table 4-26 General - Custom Properties (Inherit from Project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
--	custom	--	Custom properties can be added as needed to attach user-defined tags to projects and shares.

The following table shows **Protocols - NFS** filesystem properties.

Exceptions to the overall sharing modes may be defined for clients or collections of clients. For more information, see NFS Protocol Share Mode Exceptions.

Table 4-27 Protocols - NFS Filesystem Properties

BUI Name	CLI Name	Property Type	Description
NFS	`sharenfs`	Inherited	NFS Protocol property settings and values are described in NFS Protocol Properties.

The following table shows **Protocols - SMB** filesystem properties.

Exceptions to the overall sharing modes may be defined for clients or collections of clients. For more information, see SMB Protocol Share Mode Exceptions.

Table 4-28 Protocols - SMB Filesystem Properties

BUI Name	CLI Name	Property Type	Description
SMB	`sharesmb`	Inherited	SMB Protocol property settings and values are described in SMB Protocol Properties.

The following table shows **Protocols - Share Level ACL** filesystem properties.

Table 4-29 Protocols - Share Level ACL Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Type	--	--	Indicates the type of the ACL.
Target	--	--	Indicates the target for the ACL.
Access	--	--	Indicates whether the ACL access is allowed or denied.
Permissions: Inheritance	--	--	Specifies standard UNIX permissions for the ACL.

The following table shows **Protocols - HTTP** (Inherit from project) filesystem properties.

Table 4-30 Protocols - HTTP (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`sharedav`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Protocols - FTP** (Inherit from project) filesystem properties.

Table 4-31 Protocols - FTP (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`shareftp`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Protocols - SFTP** (Inherit from project) filesystem properties.

Table 4-32 Protocols - SFTP (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`sharesftp`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Protocols - TFTP** (Inherit from project) filesystem properties.

Table 4-33 Protocols - TFTP (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Share mode	`sharetftp`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. In the CLI, `on` is an alias for `rw`.

The following table shows **Access - File Retention Policy** (Inherit from project) filesystem properties.

Table 4-34 Access - File Retention Policy (Inherit from Project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Minimum file retention period	`retention.period.min`	Inherited	Minimum amount of time for file retention. Set a value and time measurement. Default value is `0` (zero), and the value must be less than 100 years.
Maximum file retention period	`retention.period.max`	Inherited	Maximum amount of time for file retention. Set a value and time measurement. Default value is `5 years`, and the value must be less than 100 years.
Default file retention period	`retention.period.default`	Inherited	Default amount of time for which a file is retained if it is automatically retained, or retained manually without first changing the file's access time attribute. Set a value and time measurement. Default value is `0` (zero), and the value must be between the minimum and maximum retention periods, inclusive.

Table 4-34 (Cont.) Access - File Retention Policy (Inherit from Project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Automatic file retention grace period	retention.period.grace	Inherited	Amount of time a file must remain unmodified before it is automatically retained at the default file retention period value. Set a value and time measurement. The grace period is not constrained by either the minimum period nor the maximum period.
File retention on expiry policy	retention.policy.onexpiry	Inherited	Determines behavior when file retention expires. Sets one of the following options: • Off (default): The file is not affected, and it remains on the system after retention expires. • Delete: The file is deleted after expiration has been met. • Hold: The file transitions to an indefinite hold. The file cannot be deleted until set to off or delete.
Delete expired files after	retention.period.deletegrace	Inherited	Number of seconds/hours/days/years that automatic file deletion is delayed when the file retention on expiry policy is set to delete. Default value: 0 days.
Allow permission changes on retained files	retention.policy.changeacl	Inherited	Determines if a retained file's ACL/permissions can be changed. Default value is off and the file's ACL/permissions cannot be changed. When set to on, the ACL settings/permissions, other than write, can be changed on a retained file.

The following table shows **Access - Root Directory Access** filesystem properties.

Table 4-35 Access - Root Directory Access Filesystem Properties

BUI Name	CLI Name	Property Type	Description
User	root_user	Filesystem local	Specifies the owner of the root directory.
Group	root_group	Filesystem local	Specifies the group of the root directory.
Permissions	root_permissions	Filesystem local	Specifies standard UNIX permissions for the root directory.

The following table shows **Access - ACL Behavior** (Inherit from project) filesystem properties.

Table 4-36 Access - ACL Behavior (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
ACL behavior on mode change	aclmode	Inherited	Controls how a mode change request interacts with the existing ACL.

Table 4-36   (Cont.) Access - ACL Behavior (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
ACL inheritance behavior	`aclinherit`	Inherited	Controls how a new file or directory inherits existing ACL settings from the parent directory.

The following table shows **Access - Root Directory ACL** filesystem properties.

Table 4-37   Access - Root Directory ACL Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Type	--	--	Indicates the type of the ACL.
Target	--	--	Indicates the target of the ACL.
Access	--	--	Indicates whether the ACL access is allowed or denied.
Permissions:Inheritance	--	--	Specifies standard UNIX permissions for the ACL.

The following table shows **Snapshots - Properties** (Inherit from project) filesystem properties.

Table 4-38   Snapshots - Properties (Inherit from project) Filesystem Properties

BUI Name	CLI Name	Property Type	Description
.zfs/snapshot visibility	`snapdir`	Inherited	Controls whether filesystem snapshots can be accessed over data protocols at `.zfs/snapshot` in the root of the filesystem.
Scheduled snapshot label	`snaplabel`	Inherited	Appends a user-defined label to each scheduled snapshot and is blank by default.

The following table shows **Snapshots - Snapshots** filesystem properties.

Table 4-39   Snapshots - Snapshots Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Name	--	--	Specifies the name of the snapshot.
Creation	--	--	Specifies the date and time when the snapshot is created.
Unique	--	--	Indicates the amount of unique space used by the snapshot.
Total	--	--	Indicates the total amount of space referenced by the snapshot. This represents the size of the filesystem at the time the snapshot was taken, and any snapshot can theoretically take up an amount of space equal to the total size as data blocks are rewritten.

Table 4-39 (Cont.) Snapshots - Snapshots Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Clones	--	--	Shows the number of clones of the snapshot.

The following table shows **Snapshots - Schedule** filesystem properties.

Table 4-40 Snapshots - Schedule Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Frequency	`frequency`	Create time	Indicates how often the snapshot is taken.
Keep at most	`keep`	Create time	Controls the retention policy for snapshots.

The following table shows **Replication** (Inherit from project)**/Create New Actions** filesystem properties.

Table 4-41 Replication (Inherit from project)/Create New Actions Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Target	`target`	Inherited	Identifies the replication target system.
Pool	`pool`	Inherited	Specifies the storage pool on the target where the project will be replicated.
Export data path	`export_path`	Inherited	Indicates the export data path.
Limit bandwidth	`max_bandwidth`	Inherited	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second).
Enable SSL-encryption	`use_ssl`	Inherited	Controls whether to encrypt data on the wire using SSL.
Disable compression	`compression`	Inherited	Controls whether the compression is enabled or disabled.
Include snapshot	`include_snaps`	Inherited	Controls whether replication updates include non-replication snapshots.
Retain user snapshots on target	`retain_user_snaps_on_target`	Inherited	When set, keeps user-generated snapshots on the replication target. Continues to retain snapshots on the target until disabled.
Include clone origin as data	`include_clone_origin_as_data`	Inherited	Controls the replication of each share that was cloned from a share that is external to the replication package on the target.
Recovery point objective	`recovery_point_objective`	Inherited	Specifies the maximum tolerable amount of data loss in the event of a disaster or major outage.

Table 4-41 (Cont.) Replication (Inherit from project)/Create New Actions Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Replica lag warning alert	replica_lag_warning_alert	Inherited	Specifies a limit, represented as a percentage of the RPO, when a minor alert is generated.
Replica lag error alert	replica_lag_error_alert	Inherited	Specifies a limit, represented as a percentage of the RPO, when a major alert is generated.
Update frequency	continuous	Inherited	Controls whether this action is being replicated continuously or at manual or scheduled intervals.

The following table shows **Usage** filesystem properties.

Table 4-42 Usage Filesystem Properties

BUI Name	CLI Name	Property Type	Description
Referenced data	space_data	Read-only	Shows the total amount of space referenced by the active share, independent of any snapshots.
Unused Reservation	space_unused_res	Read-only	Shows the amount of remaining space that is reserved for the filesystem.
Snapshot data	space_snapshots	Read-only	Shows the total amount of data currently held by all snapshots of the share.
Available data	space_available	Read-only	Shows any quotas on the share or project, or the absolute capacity of the pool.
Total space	space_total	Read-only	Shows the sum of referenced data, snapshot data, and unused reservation.

LUN Properties

In the CLI, use the `get` command to see a list of all properties.

The following table shows **Create LUN** properties.

Table 4-43 Create LUN Properties

BUI Name	CLI Name	Property Type	Description
Project	--	--	Defines which project the LUN uses to inherit parameter settings.
Name	--	--	Defines the name of the LUN.
Volume size	volsize	LUN local	Defines the maximum volume size and unit of measurement. For more information, see Volume size.

Table 4-43   (Cont.) Create LUN Properties

BUI Name	CLI Name	Property Type	Description
Thin provisioned	sparse	LUN local	Indicates only the amount of space physically consumed by data is used when selected. For more information, see Thin provisioned.
Space reclamation	space_reclamation	LUN local	Indicates if thin provisioned LUN free space is returned to the storage pool. For more information, see Space Reclamation.
Volume block size	volblocksize	Create time	Native block size for the LUN; any power of 2 from 512 bytes to 1M, and the default is 8K.
Online	status	LUN local	Indicates whether it is online or not.
Target group	targetgroup	LUN local	Shows groups of targets used when exporting a LUN.
Initiator group(s)	initiatorgroup	LUN local	Shows groups of initiators that can access the LUN.
Mountpoint	mountpoint	Inherited	Controls the path used to export filesystems. For more information, see Mountpoint.
LU number	lunumber	LUN local	Sets the logical unit number to 0 (zero) or automatically assigns the number.
Encryption	encryption	Inherited	Defines the encryption type. For more information see, Managing Encryption Keys.
Inherit key	--	--	If selected, indicates that the encryption key is inherited from the parent project.
Key	key	Inherited	Sets a specific encryption key that is used when the key is not inherited from the parent project.
Keyname	keyname	Static	Identifies the key.
GUID	lunguid	Read-only, LUN local	A globally unique, read-only identifier that identifies the SCSI device.

The following table shows **General - Space Usage - Data** LUN properties.

Table 4-44   General - Space Usage - Data LUN Properties

BUI Name	CLI Name	Property Type	Description
Volume size	volsize	LUN local	Defines the maximum volume size and unit of measurement. For more information, see Volume size.
Thin provisioned	sparse	LUN local	Indicates only the amount of space physically consumed by data is used when selected. For more information, see Thin provisioned.

Table 4-44 (Cont.) General - Space Usage - Data LUN Properties

BUI Name	CLI Name	Property Type	Description
Space reclamation	`space_reclamation`	LUN local	Indicates if thin provisioned LUN free space is returned to the storage pool. For more information, see Space Reclamation.

The following table shows **General - Properties** (Inherit from project) LUN properties.

Table 4-45 General - Properties (Inherit from project) LUN Properties

BUI Name	CLI Name	Property Type	Description
Data deduplication (warning)	`dedup`	Inherited	Controls whether duplicate copies of data are eliminated. For more information, see Data Deduplication.
Data compression	`compression`	Inherited	Controls whether data is compressed before being written to disk. For more information, see Data Compression.
Checksum	`checksum`	Inherited	Controls the checksum used for data blocks. For more information, see Checksum.
Additional replication	`copies`	Inherited	Controls number of copies stored of each block, above and beyond any redundancy of the storage pool. For more information, see Additional replication.
Cache device usage	`secondarycache`	Inherited	Controls whether cache devices are used for the share. For more information, see Cache device usage.
Synchronous write bias	`logbias`	Inherited	Controls the behavior when servicing synchronous writes. For more information, see Synchronous write bias.
Prevent destruction	`nodestroy`	Inherited	Prevents shares or projects from being destroyed when set. For more information, see Prevent destruction. For preventing destruction at the storage pool level, see Destroy Prevention and Approval.

The following table shows **Custom Properties** (Inherit from Project) LUN properties.

Table 4-46 Custom Properties (Inherit from Project) LUN Properties

BUI Name	CLI Name	Property Type	Description
Schema	`schema`	--	Custom properties can be added as needed to attach user-defined tags to projects and shares. For more information, see Working with Schemas.

The following table shows **Protocols - Sharing Options** LUN properties.

Table 4-47 Protocols - Sharing Options LUN Properties

BUI Name	CLI Name	Property Type	Description
Online	`status`	LUN local	Indicates whether it is online or not.
Target group	`targetgroup`	LUN local	Shows groups of targets used when exporting a LUN.
Initiator group(s): LU number	`initiatorgroup`	LUN local	Shows groups of initiators that can access the LUN.

The following table shows **Protocols - Write Cache Behavior** LUN properties.

Table 4-48 Protocols - Write Cache Behavior LUN Properties

BUI Name	CLI Name	Property Type	Description
Write cache enabled	`writecache`	LUN local	Controls whether the LUN caches writes.

The following table shows **Snapshots - Properties** (Inherit from project) LUN properties.

Table 4-49 Snapshots - Properties (Inherit from project) LUN Properties

BUI Name	CLI Name	Property Type	Description
Scheduled snapshot label	`snaplabel`	Inherited	Appends a user-defined label to each scheduled snapshot and is blank by default.

The following table shows **Snapshots - Snapshots** LUN properties.

Table 4-50 Snapshots - Snapshots LUN Properties

BUI Name	CLI Name	Property Type	Description
Name	--	--	Specifies the name of the snapshot.
Creation	--	--	Specifies the date and time when the snapshot is created.
Unique	--	--	Indicates the amount of unique space used by the snapshot.
Total	--	--	Indicates the total amount of space referenced by the snapshot.
Clones	--	--	Shows the number of clones of the snapshot.

The following table shows **Snapshots - Schedules** LUN properties.

Table 4-51 Snapshots - Schedules LUN Properties

BUI Name	CLI Name	Property Type	Description
Frequency	`frequency`	Create time	Indicates how often the snapshot is taken.

Table 4-51   (Cont.) Snapshots - Schedules LUN Properties

BUI Name	CLI Name	Property Type	Description
Keep at most	`keep`	Create time	Controls the retention policy for snapshots.

The following table shows **Replication** (Inherit from project)**/Create New Actions** LUN properties.

Table 4-52   Replication (Inherit from project)/Create New Actions LUN Properties

BUI Name	CLI Name	Property Type	Description
Target	`target`	Inherited	Identifies the replication target system.
Pool	`pool`	Inherited	Specifies the storage pool on the target where the project will be replicated.
Export data path	`export_path`	Inherited	Indicates the export data path.
Limit bandwidth	`max_bandwidth`	Inherited	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second).
Enable SSL-encryption	`use_ssl`	Inherited	Controls whether to encrypt data on the wire using SSL.
Disable compression	`compression`	Inherited	Controls whether the compression is enabled or disabled.
Include snapshot	`include_snaps`	Inherited	Controls whether replication updates include non-replication snapshots.
Retain user snapshots on target	`retain_user_snaps_ on_target`	Inherited	When set, keeps user-generated snapshots on the replication target.
Update frequency	`continuous`	Inherited	Controls whether this action is being replicated continuously or at manual or scheduled intervals.

The following table shows **Usage** LUN properties.

Table 4-53   Usage LUN Properties

BUI Name	CLI Name	Property Type	Description
Referenced data	`space_data`	Read-only	Shows the total amount of space referenced by the active share, independent of any snapshots.
Snapshot data	`space_snapshots`	Read-only	Shows the total amount of data currently held by all snapshots of the share.
Available data	`space_available`	Read-only	Shows any quotas on the share or project, or the absolute capacity of the pool.
Total space	`space_total`	Read-only	Shows the sum of referenced data, snapshot data, and unused reservation.

Space Management for Shares

Oracle ZFS Storage Appliance manages physical storage using a pooled storage model where all filesystems and LUNs share common space. Filesystems never have an explicit size assigned to them, and only take up as much space as they need. LUNs reserve enough physical space to write the entire contents of the device, unless they are thin provisioned, in which case they behave like filesystems and use only the amount of space physically consumed by data.

In addition, thin provisioned LUNs can optionally take advantage of the space reclamation feature that returns free space to the storage pool, as described in Space Reclamation. The logical block provisioning (LBP) threshold for thin provisioned LUNs is another way to manage space. The `LBPthreshold` command is available in the RESTful API only, and it sets the LBP threshold for thin provisioned LUNs within a specified storage pool. When set to the default value of 0, the LBP threshold is disabled. To set the threshold, specify the average storage consumption rate. When the threshold value is exceeded, an error message is returned, and you can manage the space accordingly. See also Set LBP Threshold in *Oracle ZFS Storage Appliance RESTful API Guide, Release OS8.8.x*.

This system provides maximum flexibility and simplicity of management in an environment when users are generally trusted to do the right thing. A stricter environment, where a user's data usage is monitored and/or restricted, requires more careful management.

These topics define terminology and how to manage space usage, on a per-share or per-user basis, using quotas and reservations.

- Shares Terminology
- Managing Filesystem and Project Space
- Setting User or Group Quotas
- Working with Identity Management
- Share Usage Statistics

Shares Terminology

Before getting into details, it is important to understand some basic terms used when talking about space usage on Oracle ZFS Storage Appliance.

- **Physical Data** - Size of data as stored physically on disk. Typically, this is equivalent to the logical size of the corresponding data, but can be different in the phase of compression or other factors. This includes the space of the active share as well as all snapshots. Space accounting is generally enforced and managed based on physical space.
- **Logical Data** - The amount of space logically consumed by a filesystem. This does not factor into compression, and can be viewed as the theoretical upper bound on the amount of space consumed by the filesystem. Copying the filesystem to another appliance using a different compression algorithm will not consume more than this amount. This statistic is not explicitly exported and can generally only be computed by taking the amount of physical space consumed and multiplying by the current compression ratio.
- **Referenced Data** - This represents the total amount of space referenced by the active share, independent of any snapshots. This is the amount of space that the share would consume should all snapshots be destroyed. This is also the amount of data that is directly manageable by the user over the data protocols.

- **Snapshot Data** - This represents the total amount of data currently held by all snapshots of the share. This is the amount of space that would be free should all snapshots be destroyed.
- **Quota** - A quota represents a limit on the amount of space that can be consumed by any particular entity. This can be based on filesystem, project, user, or group, and is independent of any current space usage.
- **Reservation** - A reservation represents a guarantee of space for a particular project or filesystem. This takes available space away from the rest of the pool without increasing the actual space consumed by the filesystem. This setting cannot be applied to users and groups. The traditional notion of a statically sized filesystem can be created by setting a quota and reservation to the same value.

Managing Filesystem and Project Space

The simplest way of enforcing quotas and reservations is on a per-project or per-filesystem basis. Quotas and reservations do not apply to LUNs, though their usage is accounted for in the total project quota or reservations.

Data Quotas - A data quota enforces a limit on the amount of space a filesystem or project can use. By default, it will include the data in the filesystem and all snapshots. Clients attempting to write new data will get an error when the filesystem is full, either because of a quota or because the storage pool is out of space. As described in Snapshot Space Management, this behavior may not be intuitive in all situations, particularly when snapshots are present. Removing a file may cause the filesystem to write new data if the data blocks are referenced by a snapshot, so it may be the case that the only way to decrease space usage is to destroy existing snapshots.

If the "include snapshots" property is unset, then the quota applies only to the immediate data referenced by the filesystem, not any snapshots. The space used by snapshots is enforced by the project-level quota but is otherwise not enforced. In this situation, removing a file referenced by a snapshot will cause the filesystem's referenced data to decrease, even though the system as a whole is using more space. If the storage pool is full (as opposed to the filesystem reaching a preset quota), then the only way to free up space may be to destroy snapshots.

Data quotas are strictly enforced, which means that as space usage nears the limit, the amount of data that can be written must be throttled as the precise amount of data to be written is not known until after writes have been acknowledged. This can affect performance when operating at or near the quota. Because of this, it is generally advisable to remain below the quota during normal operating procedures.

Quotas are managed through the BUI under **Shares: General: Space Usage: Data**. They are managed in the CLI as the `quota` and `quota_snap` properties.

Data Reservations - A data reservation is used to make sure that a filesystem or project has at least a certain amount of available space, even if other shares in the system try to use more space. This unused reservation is considered part of the filesystem, so if the rest of the pool (or project) reaches capacity, the filesystem can still write new data even though other shares may be out of space.

By default, a reservation includes all snapshots of a filesystem. If the `include snapshots` property is unset, then the reservation only applies to the immediate data of the filesystem. The behavior when taking snapshots may not always be intuitive. If a reservation on filesystem data (but not snapshots) is in effect, then whenever a snapshot is taken, the system must reserve enough space for that snapshot to diverge completely, even if that never occurs. For example, if a 50G filesystem has a 100G reservation without snapshots, then taking the first snapshot will reserve an additional 50G of space, and the filesystem will end up reserving 150G of space total. If there is insufficient space to guarantee complete divergence of data, then taking the snapshot will fail.

Reservations are managed through the BUI under **Shares: General: Space Usage: Data**. They are managed in the CLI as the `reservation` and `reservation_snap` properties.

Space Management for Replicating LUNs - When you create a LUN the full physical space you configure for the LUN is reserved and cannot be used by other file systems (unless it is thinly provisioned). For replication, when you take a snapshot of a LUN of any given size, up to twice the size of the LUN is also reserved, depending on how much of the LUN space has been used.

The following list shows the maximum overhead space required when replicating a LUN:

- Up to 100% on the source between updates
- Up to 200% on the source during an update
- Up to 200% on the target

Setting User or Group Quotas

Quotas can be set on a user or group at the filesystem level, as well as the project level. These enforce physical data usage based on the POSIX or Windows identity of the owner or group of the file or directory. There are some significant differences between user and group quotas, and filesystem and project data quotas:

- User and group quotas can be applied to filesystems and projects.
- Default quotas can be set at the project level and inherited by the project's filesystems.
- Default quotas set at the project level can be changed at the filesystem level.
- Default quotas can be retrieved or modified over the SMB protocol.
- User and group quotas are implemented using *delayed enforcement*. This means that users will be able to exceed their quota for a short period of time before data is written to disk. After the data has been pushed to disk, the user will receive an error on new writes, just as with the filesystem-level quota case.
- User and group quotas are always enforced against referenced data. This means that snapshots do not affect any quotas, and a clone of a snapshot will consume the same amount of effective quota, even though the underlying blocks are shared.
- User and group reservations are not supported.
- User and group quotas, unlike data quotas, are stored with the regular filesystem data. This means that if the filesystem is out of space, you will not be able to make changes to user and group quotas. You must first make additional space available before modifying user and group quotas.
- User and group quotas are sent as part of any remote replication. It is up to the administrator to ensure that the name service environments are identical on the source and destination.
- An NDMP backup-and-restore operation of an entire share will include any user or group quotas. Restores into an existing share will not affect any current quotas.

Working with Identity Management

User and group quotas leverage the identity mapping service on Oracle ZFS Storage Appliance. This allows users and groups to be specified as either UNIX or Windows identities, depending on the environment. Like file ownership, these identities are tracked in the following ways:

- If there is no UNIX mapping, a reference to the windows ID is stored.
- If there is a UNIX mapping, then the UNIX ID is stored.

This means that the canonical form of the identity is the UNIX ID. If the mapping is changed later, the new mapping will be enforced based on the new UNIX ID. If a file is created by a Windows user when no mapping exists, and a mapping is later created, new files will be treated as a different owner for the purposes of access control and usage format. This also implies that if a user ID is reused (that is, a new user name association created), then any existing files or quotas will appear to be owned by the new user name.

It is recommended that any identity mapping rules be established before attempting to actively use filesystems. Otherwise, any change in mapping can sometimes have surprising results.

Working with Filesystem Namespace

Every filesystem on Oracle ZFS Storage Appliance must be given a unique mountpoint which serves as the access point for the filesystem data. Projects can be given mountpoints, but these serve only as a tool to manage the namespace using inherited properties. Projects are never mounted, and do not export data over any protocol.

All shares must be mounted under `/export`. While it is possible to create a filesystem mounted at `/export`, it is not required. If such a share does not exist, any directories will be created dynamically as necessary underneath this portion of the hierarchy. Each mountpoint must be unique within a cluster.

- **Namespace Nested Mountpoints -** It is possible to create filesystems with mountpoints beneath that of other filesystems. In this scenario, the parent filesystems are mounted before children and vice versa. The following cases should be considered when using nested mountpoints:
  - If the mountpoint does not exist, one will be created, owned by root and mode `0755`. This mountpoint may or may not be torn down when the filesystem is renamed, destroyed, or moved, depending on circumstances. To be safe, mountpoints should be created within the parent share before creating the child filesystem.
  - If the parent directory is read-only, and the mountpoint does not exist, the filesystem mount will fail. This can happen synchronously when creating a filesystem, but can also happen asynchronously when making a large-scale change, such as renaming filesystems with inherited mountpoints.
  - When renaming a filesystem or changing its mountpoint, all children beneath the current mountpoint as well as the new mountpoint (if different) will be unmounted and remounted after applying the change. This will interrupt any data services currently accessing the share.
  - Support for automatically traversing nested mountpoints depends on protocol, as outlined in the following.
- **Namespace NFSv2 / NFSv3 / NFSv4.0 / NFSv4.1 -** Under NFS, each filesystem is a unique export made visible via the `MOUNT` protocol. NFSv2 and NFSv3 have no way to

traverse nested filesystems, and each filesystem must be accessed by its full path. While nested mountpoints are still functional, attempts to cross a nested mountpoint will result in an empty directory on the client. While this can be mitigated through the use of automount mounts, transparent support of nested mountpoints in a dynamic environment requires NFSv4.0 or NFSv4.1.

NFSv4.0 and NFSv4.1 have several improvements over NFSv3 when managing mountpoints. First is that parent directories can be mounted, even if there is no share available at that point in the hierarchy. For example, if `/export/home` was shared, it is possible to mount `/export` on the client and traverse into the actual exports transparently. More significantly, some NFSv4.0 and NFSv4.1 clients (including Linux) support automatic client-side mounts, sometimes referred to as "mirror mounts." With such a client, when a user traverses a mountpoint, the child filesystem is automatically mounted at the appropriate local mountpoint, and torn down when the filesystem is unmounted on the client. From the server's perspective, these are separate mount requests, but they are stitched together onto the client to form a seamless filesystem namespace.

- **Namespace SMB** - The SMB protocol does not use mountpoints, as each share is made available by resource name. However, each filesystem must still have a unique mountpoint. Nested mountpoints (multiple filesystems within one resource) are not currently supported, and any attempt to traverse a mountpoint will result in an empty directory.
- **Namespace FTP / FTPS / SFTP** - Filesystems are exported using their standard mountpoint. Nested mountpoints are fully supported and are transparent to the user. However, it is not possible to not share a nested filesystem when its parent is shared. If a parent mountpoint is shared, then all children will be shared as well.
- **Namespace HTTP / HTTPS** - Filesystems are exported under the `/shares` directory, so a filesystem at `/export/home` will appear at `/shares/export/home` over HTTP/HTTPS. Nested mountpoints are fully supported and are transparent to the user. The same behavior regarding conflicting share options described in the FTP protocol section also applies to HTTP.

Share Usage Statistics

On the left side of the view (beneath the **Project** panel when collapsed) is a table explaining the current space usage statistics. These statistics are either for a particular share (when editing a share) or for the storage pool as a whole (when looking at the list of shares). If any properties are `0` (zero), then they are excluded from the table.

Some of the usage statistics are also displayed in the CLI context `shares show`.

The following table describes the BUI and CLI usage properties.

BUI Name	CLI Name	Description
Available space	--	This statistic is implicitly shown as the capacity in terms of capacity percentage in the title. The available space reflects any quotas on the share or project, or the absolute capacity of the pool. The number shown here is the sum of the total space used and the amount of available space.

BUI Name	CLI Name	Description
Referenced data	`usage_data`	The amount of data referenced by the data. This includes all filesystem data or LUN blocks, in addition to requisite metadata. With compression, this value may be much less than the logical size of the data contained within the share. If the share is a clone of a snapshot, this value may be less than the physical storage it could theoretically include, and may be `0` (zero).
Snapshot data	`usage_snapshots`	The amount of space used by all snapshots of the share, including any project snapshots. This size is not equal to the sum of unique space consumed by all snapshots. Blocks that are referenced by multiple snapshots are not included in the per-snapshot usage statistics, but will show up in the share's snapshot data total.
Unused reservation	`--`	If a filesystem has a reservation set, this value indicates the amount of remaining space that is reserved for the filesystem. This value is not set for LUNs. The appliance prevents other shares from consuming this space, guaranteeing the filesystem enough space. If the reservation does not include snapshots, then there must be enough space when taking a snapshot for the entire snapshot to be overwritten. For more information on reservations, see Filesystem Properties.
Total space	`usage_total`	The sum of referenced data, snapshot data, and unused reservation.

Share and Project Protocols

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project.

For iSCSI, initiators can discover the target through one of the mechanisms described in Configuring Storage Area Network (SAN).

For information about supported protocol properties, see the following sections:

- NFS Protocol
- SMB Protocol
- HTTP Protocol
- FTP Protocol
- SFTP Protocol
- TFTP Protocol

Related Topics

- NFS Configuration
- SMB Configuration

NFS Protocol

This section contains the following topics:

- NFS Protocol Properties

- NFS Protocol Share Mode Exceptions
- NFS Protocol Character Set Encodings
- NFS Protocol Security Modes

For more information about the NFS protocol, use these topics:

- NFS Configuration
- Filesystem Properties
- Project Properties
- NFSv2 and NFSv3 Security (RFC 2623) (http://www.ietf.org/rfc/rfc2623.txt)
- NFSv4 Protocol (RFC 7530) (http://www.ietf.org/rfc/rfc7530.txt)
- NFSv4.1 Protocol (RFC 5661) (https://tools.ietf.org/html/rfc5661)

For information about other supported protocols, see the following sections:

- SMB Protocol
- HTTP Protocol
- FTP Protocol
- SFTP Protocol
- TFTP Protocol

NFS Protocol Properties

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. The following table shows NFS protocol properties and possible values.

Table 4-54 NFS Protocol Properties

Property	CLI Value(s)	Property Type	Description
Share mode	`off` | `rw` | `ro`	Inherited	Determines whether the share is available for reading only, for reading and writing, or neither. See Share and Project Protocols.
Disable setuid/ setgid file creation	`nosuid`	Inherited	If selected, clients will not be able to create files with the `setuid (S_ISUID)` and `setgid (S_ISGID)` bits set, nor enable these bits on existing files via the `chmod(2)` system call.
Prevent clients from mounting subdirectories	`nosub`	Inherited	If selected, clients will be prevented from directly mounting subdirectories. They will be forced to mount the root of the share. Note: This only applies to the NFSv2 and NFSv3 protocols, not to NFSv4.0 or NFSv4.1.

Table 4-54 (Cont.) NFS Protocol Properties

Property	CLI Value(s)	Property Type	Description
Anonymous user mapping	`anon`	Inherited	Unless the `root` option is in effect for a particular client, the root user on that client is treated as an unknown user, and all attempts by that user to access the sharer's files will be treated as attempts by a user with this UID. This file's access bits and ACLs will then be evaluated normally.
Character set	See Character Set Encodings for possible values.	Inherited	Sets the character set default for all clients.
Security mode	`sec=` See Security Modes for list of possible values.	Inherited	Sets the security mode for all clients.
Enforce reserved ports for system authentication	`resvport`	Inherited	When set on a share or project in conjunction with the system authentication security mode, requires NFS clients to use low-numbered ("reserved") TCP ports. Some NFS clients, such as Oracle Solaris and Linux, use low-numbered TCP ports by default. Other clients, such as Microsoft Windows, may require configuration.

NFS Protocol Share Mode Exceptions

Exceptions to the global sharing mode may be defined for clients or collections of clients by setting client-specific share modes or *exceptions*. To restrict access to certain clients, set the global sharing mode to `none`, and increasingly grant access to smaller and smaller groups. For example, you could create a share with the global sharing mode set to `none`, which denies access to all clients, and then grant read-only access to a subset of the clients. Further, you could grant read/write access to an even smaller subset of the clients and, finally, only trusted hosts might have read/write and root-enabled access.

Client-specific share modes take precedence over the global share mode. A client is granted access according to the client-specific share mode that is specified in an exception. In the absence of exceptions, the client is granted access according to the global share mode.

Table 4-55 Client Types

Type	CLI Prefix	Description	Example
Host (FQDN)	`none`	A single client with an IP address that resolves to the specified fully qualified name.	`hostname.sf.example.com`

Table 4-55 (Cont.) Client Types

Type	CLI Prefix	Description	Example
Netgroup	%	A netgroup name in LDAP that grants access to certain named clients. This client type can only be used in an exception if the `explicit_netgroups` property is set to `true` in the CLI, or **Use new syntax for netgroups in share access lists** is selected in the BUI.	`netgroup.sf.example.com`
DNS Domain	.	All clients with IP addresses that resolve to a fully qualified name ending in this suffix.	`sf.example.com`
IPv4 Subnet	@	All clients with IP addresses that are within the specified IPv4 subnet, expressed in CIDR notation.	`192.0.2.254/22`
IPv6 Subnet	@	All clients with IP addresses that are within the specified IPv6 subnet, expressed in CIDR notation.	`2001:db8:410:d43::/64`

For each client or collection of clients, you specify whether the client has read-only or read-write access to the share. If you are setting an NFS exception, you also specify whether the client has root user privileges or is treated as a user without root access.

Managing Netgroups

Netgroups can be used to control access for NFS exports. However, managing netgroups can be complex. Consider using IP subnet rules or DNS domain rules instead.

If netgroups are used, they will be resolved from NIS or LDAP, depending on which service is enabled. If LDAP is used, each netgroup must be located at the default location, `ou=Netgroup`, (Base DN), and must use the standard schema.

The username component of a netgroup entry typically has no effect on NFS; only the hostname is significant. Hostnames contained in netgroups must be canonical and, if resolved using DNS, fully qualified. That is, the NFS subsystem will attempt to verify that the IP address of the requesting client resolves to a canonical hostname that matches either the specified FQDN, or one of the members of one of the specified netgroups. This match must be exact, including any domain components; otherwise, the exception will not match and the next exception will be tried. For more information on hostname resolution, see DNS.

As of the 2013.1.0 software release, UNIX client users may belong to a maximum of 1024 groups without any performance degradation. Prior releases supported up to 16 groups per UNIX client user.

NFS Share Modes and Exception Options

In the CLI, all NFS share modes and exceptions are specified using a single options string for the `sharenfs` property. This string is a comma-separated list of values. It should begin with one of `ro`, `rw`, `on`, or `off`, as an analogue to the global share modes described for the BUI.

Table 4-56 NFS Share Mode Values (BUI and CLI)

BUI Share Mode Value	CLI Share Mode Value	Description	Example
None	`off`	Share mode is disabled.	`sharenfs=off`
	`on`	The share name is the dataset name and is available for reading and writing or reading only if the `rw` or `ro` NFS exceptions are defined. For all other clients, share mode is disabled.	`sharenfs="on,ro=sf.example.com"`
	`resource_name`	The share name is the resource name and is available for reading and writing or reading only if the `rw` or `ro` NFS exceptions are defined. For all other clients, share mode is disabled.	`sharenfs="myshare,ro=sf.example.com"`
Read/write	`on`	The share name is the dataset name and is available for reading and writing for all clients if there are no NFS exceptions.	`sharenfs=on`
	`rw`	The share name is the dataset name and is available for reading and writing for all clients except those for which the `ro` exception is defined.	`sharenfs=rw` or `sharenfs="rw,ro=sf.example.com"`
	`resource_name`	The share name is the resource name and is available for reading and writing for all clients if there are no NFS exceptions.	`sharenfs=myshare`
	`resource_name``,rw`	The share name is the resource name and is available for reading and writing for all clients except those for which the `ro` exception is defined. NFS exceptions may or may not be defined.	`sharenfs="myshare,rw"` or `sharenfs="myshare,rw,ro=sf.example.com"`
Read only	`ro`	The share name is the dataset name and is available for reading only for all hosts except those for which the `rw` exception is defined.	`sharenfs="ro,rw=sf.example.com"`
	`resource_name``,ro`	The share name is the resource name and is available for reading only for all clients except those for which the `rw` exception is defined. NFS exceptions may or may not be defined.	`sharenfs="myshare,ro"` or `sharenfs="myshare,ro,rw=sf.example.com"`

The following example sets the share mode for all clients to read-only. The root users on all clients will access the files on the share as if they were the generic `nobody` user.

```
set sharenfs=ro
```

Either or both of the `nosuid` and `anon` options can also be appended. Therefore, to define the mapping of all unknown users to the UID 153762, you might specify the following:

```
set sharenfs="ro,anon=153762"
```

> **Note:**
>
> CLI property values that contain the = character must be quoted.

Additional NFS exceptions can be specified by appending text of the form *`option=collection`*, where *option* is one of `ro`, `rw`, or `root`, defining the type of access to be granted to the client collection. The collection is specified by the prefix character from Client Types table and either a DNS hostname/domain name or CIDR network number. For example, to grant read-write access to all hosts in the `sf.example.com` domain and root access to those in the 192.168.44.0/24 network, you might use:

```
set sharenfs="ro,anon=153762,rw=.sf.example.com,root=@192.168.44.0/24"
```

> **Note:**
>
> This example only applies to NFS exceptions.

Netgroup names can be used anywhere an individual fully qualified hostname can be used. For example, you can permit read-write access to the `engineering` netgroup as follows:

```
set sharenfs="ro,rw=engineering"
```

NFS Protocol Character Set Encodings

Normally, the character set encoding used for filename is unspecified. The NFSv3 and NFSv2 protocols do not specify the character set. NFSv4.0 and NFSv4.1 are supposed to use UTF-8, but not all clients do and this restriction is not enforced by the server. If the UTF-8 only option is disabled for a share, these filenames are written verbatim to the filesystem without any knowledge of their encoding. This means that they can only be interpreted by clients using the same encoding. SMB, however, requires filenames to be stored as UTF-8 so that they can be interpreted on the server side. This makes it impossible to support arbitrary client encodings while still permitting access over SMB.

To support such configurations, the character set encoding can be set share-wide or on a per-client basis. The following character set encodings are supported:

cp932	euc-tw	iso8859-7	koi8-r
euc-cn	iso8859-1	iso8859-8	shift_jis
euc-jp	iso8859-2	iso8859-9	
euc-jpms	iso8859-5	iso8859-13	
euc-kr	iso8859-6	iso8859-15	

The default behavior is to leave the character set encoding unspecified (pass-through). The BUI allows the character set to be chosen through the standard exception list mechanism. In the CLI, each character set itself becomes an option with one or more hosts, with `*` indicating the share-wide setting. For example, the following:

```
hostname:shares default> set sharenfs="rw,euc-kr=*"
```

Will share the filesystem with `euc-kr` as the default encoding. The following:

```
hostname:shares default> set sharenfs="rw,euc-kr=host1.example.com,euc-
jp=host2.example.com"
```

Use the default encoding for all clients except `host1` and `host2`, which will use `euc-kr` and `euc-jp`, respectively. The format of the host lists follows that of other CLI NFS options.

Note that some NFS clients do not correctly support alternate locales; consult your NFS client documentation for details.

NFS Protocol Security Modes

Security modes are set on a per-share basis. The following list describes Kerberos security settings:

- **krb** - End-user authentication through Kerberos V5
- **krb5i** - krb5 plus integrity protection (data packets are tamper proof
- **krb5p** - krb5i plus privacy protection (data packets are tamper proof and encrypted)

Security modes are specified by appending text in the form *`option=mode`* where *option* is `sec` and *mode* is the security setting. For example:

```
hostname: shares default> set sharenfs="sec=krb5"
```

> **Note:**
>
> CLI property values that contain the `=` character must be quoted.

Combinations of Kerberos types can be specified in the security mode setting. The combination security modes let clients mount with any Kerberos type listed, as shown in the following table.

Table 4-57 Combinations of Kerberos types

Setting	Description
`sys`	System authentication.
`krb5`	Kerberos v5 only - Clients must mount using this flavor.
`krb5:krb5i`	Kerberos v5, with integrity - Clients may mount using any flavor listed.
`krb5i`	Kerberos v5 integrity only - Clients must mount using this flavor.
`krb5:krb5i:krb5p`	Kerberos v5, with integrity or privacy - Clients may mount using any flavor listed.
`krb5p`	Kerberos v5 privacy only - Clients may mount using this flavor.

Reserved Ports

To set reserved ports for system authentication, use `resvport` as shown in this example:

```
set sharenfs="sec=sys,rw,resvport"
```

Note that `resvport` can only be used with the system authentication security mode `sec=sys`.

SMB Protocol

This section contains the following topics:

- SMB Protocol Properties
- Client-side Caching Property
- Opportunistic Locks Property
- SMB Protocol Share Mode Exceptions
- Share-Level ACLs

For more information about the SMB protocol, use these topics:

- SMB Configuration
- Filesystem Properties
- Project Properties

For information about other supported protocols, see the following sections:

- NFS Protocol
- HTTP Protocol
- FTP Protocol
- SFTP Protocol
- TFTP Protocol

SMB Protocol Properties

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project.

Table 4-58 SMB Protocol Properties

BUI Property	CLI Property	Property Type	Description
Share mode	`off` | `rw` | `ro`	Inherited	Specifies whether the share is available for reading only, for reading and writing, or neither. See table "SMB Share Mode Values (BUI and CLI)" in SMB Protocol Share Mode Exceptions.
Resource name	*`resource_name`*	Inherited	The name by which SMB clients refer to this share. Share mode exceptions can be specified for this resource. See table "SMB Share Mode Values (BUI and CLI)" in SMB Protocol Share Mode Exceptions.

Table 4-58 (Cont.) SMB Protocol Properties

BUI Property	CLI Property	Property Type	Description
Enable access-based enumeration	`abe`	Inherited	Specifies whether to perform access-based enumeration.
Enable guest access	`guestok`	Inherited	Specifies whether to grant guest access. This property is disabled by default.
Is a DFS namespace	`dfsroot`	Inherited	Specifies whether this share is provisioned as a standalone DFS namespace.
Client-side caching policy	`csc`	Inherited	Per-share configuration options provided to support client-side caching. For more information, see Client-side Caching Property.
Opportunistic locks policy	`oplocks`	Inherited	Specifies whether opportunistic locks are enabled at the share level. For more information, see Opportunistic Locks Property.
Enable continuous availability	`cont_avail`	Inherited	Specifies whether SMB3 clients can request persistent file handles for the share. When enabled, the appliance can store the state associated with a persistent file handle in stable, persistent storage. The state can be transparently restored in the event of a controller failure, such as a takeover and failback operation on clustered controllers. Continuously available SMB shares are not allowed to be shared over NFS or used on workloads such as Home Directory that have a very high number of opens/closes. Continuously available SMB shares are only recommended for enterprise applications that have limited number of opens/closes.
Encrypt data access	`encrypt`	Inherited	Specifies whether SMB3 encryption is enabled at the share level. When enabled, the SMB server requires clients to encrypt requests to access the share. This enforcement can be bypassed if the server allows unencrypted access. This property is disabled by default. For global-level SMB encryption properties, see `Encrypt data access` and `Reject unencrypted access` in SMB Service Properties.

Table 4-58 (Cont.) SMB Protocol Properties

BUI Property	CLI Property	Property Type	Description
Bypass traverse checking	`bypasstraverse`	Inherited	Specifies whether to bypass traverse checking for the share. This property is disabled by default. When bypass traverse is disabled, UNIX semantics are used: Always enforce the traversal permissions of folders when navigating an object on this share. When bypass traverse is enabled, Windows semantics are used: Access to an object on this share depends on the user's rights to that object, ignoring the traversal permissions of folders.

Client-side Caching Property

The client-side caching property (`csc`) controls whether files and programs from the share are cached on the local client for offline use when disconnected from the appliance.

BUI Value	CLI Value	Description
No caching	`none`	Disables client-side caching for the share. No files or programs from the share are available offline. This option blocks offline files on the client computers from making copies of the files and programs on the shared folder.
Manual caching	`manual`	Only specified files and programs are cached on the local client and available offline. This is the default option when you set up a shared folder. By using this option, no files or programs are available offline by default. You can control which files and programs to access when you are not connected to the network.
Automatic document caching	`documents`	All files accessed from the share are cached on the local client and available offline. Files are automatically reintegrated when the local client is online again. Programs accessed from the share are not available offline unless previously cached locally.
Automatic program caching	`programs`	All programs accessed from the share are cached on the local client and available offline. When online, the programs are run from the local client. Additionally, all files accessed from the share are cached on the local client and available offline. Files are automatically reintegrated when the local client is online again.

Opportunistic Locks Property

Opportunistic locks are a client-caching mechanism that facilitates local caching to reduce network traffic and improve performance. The property (`oplocks`) controls whether the server grants or denies opportunistic locks at the share level, and applies to both lease (SMB 2.1 and above) and legacy (SMB 2.0 and below) opportunistic locks.

The client requests an opportunistic lock on a file within a share, and that request is either granted or denied depending on the server configuration and the current state of the file. If the client attempts to access a file in a manner inconsistent with the opportunistic locks that have already been granted for that file, a conflict occurs. In such cases, the server initiates a process to break the existing opportunistic locks before proceeding with the conflicting operation.

Enabling opportunistic locks improves performance when files within a share are accessed by a single client. In some scenarios, however, such as when the same file is accessed simultaneously by multiple clients, it can introduce unnecessary overhead. Opportunistic locks can thus be enabled or disabled per share, instead of globally controlled, based on the expected pattern of workloads.

If an opportunistic locks property is not defined at the share level, the default is the global opportunistic locks property set at the service level. For more information, see `Enable oplocks` in section SMB Service Properties.

BUI Value	CLI Value	Description	Example
Enabled	`enabled`	Enables opportunistic locks for a share.	`set sharesmb="myshare,oplocks=enabled,abe=off,dfsroot=false"`
Disabled	`disabled`	Disables opportunistic locks for a share.	`set sharesmb="myshare,oplocks=disabled,abe=off,dfsroot=false"`
empty	--	The opportunistic locks property is neither enabled or disabled. Uses the global opportunistic locks property when the share-level property is not set.	`set sharesmb="myshare,abe=off,dfsroot=false"`

SMB Protocol Share Mode Exceptions

Exceptions to the global sharing mode may be defined for clients or collections of clients by setting client-specific share modes or *exceptions*. To restrict access to certain clients, set the global sharing mode to `none` and increasingly grant access to smaller and smaller groups. For example, you could create a share with the global sharing mode set to none, which denies access to all clients, and then grant read-only access to a subset of the clients. Further, you could grant read/write access to an even smaller subset of the clients and, finally, only trusted hosts might have read/write access.

Table 4-59 Client Types

Type	CLI Prefix	Description	Example
Host (FQDN)	`none`	A single client with an IP address that resolves to the specified fully qualified name.	`hostname.sf.example.com`

Table 4-59 (Cont.) Client Types

Type	CLI Prefix	Description	Example
Netgroup	%	A netgroup name in LDAP that grants access to certain named clients. This client type can only be used in an exception if the `explicit_netgroups` property is set to `true` in the CLI, or **Use new syntax for netgroups in share access lists** is selected in the BUI.	`netgroup.sf.example.com`
DNS Domain	.	All clients with IP addresses that resolve to a fully qualified name ending in this suffix.	`sf.example.com`
IPv4 Subnet	@	All clients with IP addresses that are within the specified IPv4 subnet, expressed in CIDR notation.	`192.0.2.254/22`
IPv6 Subnet	@	All clients with IP addresses that are within the specified IPv6 subnet, expressed in CIDR notation.	`2001:db8:410:d43::/64`

For each client or collection of clients, you specify whether the client has read-only or read-write access to the share.

Managing netgroups - Netgroups can be used to control access for SMB exports. However, managing netgroups can be complex. Consider using IP subnet rules or DNS domain rules instead.

If netgroups are used, they will be resolved from NIS or LDAP, depending on which service is enabled. If LDAP is used, each netgroup must be located at the default location, `ou=Netgroup`, (Base DN), and must use the standard schema.

The username component of a netgroup entry typically has no effect on SMB; only the hostname is significant. Hostnames contained in netgroups must be canonical and, if resolved using DNS, fully qualified. That is, the SMB subsystem will attempt to verify that the IP address of the requesting client resolves to a canonical hostname that matches either the specified FQDN, or one of the members of one of the specified netgroups. This match must be exact, including any domain components; otherwise, the exception will not match and the next exception will be tried. For more information on hostname resolution, see DNS.

As of the 2013.1.0 software release, UNIX client users may belong to a maximum of 1024 groups without any performance degradation. Prior releases supported up to 16 groups per UNIX client user.

SMB Share Modes and Exception Options

In the CLI, all SMB share modes and exceptions are specified using a single options string for the `sharesmb` property. This string is a comma-separated list of values. It should begin with one of `ro`, `rw`, `on`, or `off`, as an analogue to the global share modes described for the BUI.

Table 4-60 SMB Share Mode Values (BUI and CLI)

BUI Share Mode Value	CLI Share Mode Value	Description	Example
None	`off`	Share mode is disabled.	`sharesmb=off`

Table 4-60 (Cont.) SMB Share Mode Values (BUI and CLI)

BUI Share Mode Value	CLI Share Mode Value	Description	Example
	`on`	The share name is the dataset name and is available for reading and writing or reading only if the `rw` or `ro` SMB exceptions are defined. For all other clients, share mode is disabled.	`sharesmb="on,ro=sf.example.com"`
	`resource_name`	The share name is the resource name and is available for reading and writing or reading only if the `rw` or `ro` SMB exceptions are defined. For all other clients, share mode is disabled.	`sharesmb="myshare,ro=sf.example.com"`
Read/write	`on`	The share name is the dataset name and is available for reading and writing for all clients if there are no SMB exceptions.	`sharesmb=on`
	`rw`	The share name is the dataset name and is available for reading and writing for all clients except those for which the `ro` exception is defined.	`sharesmb=rw` or `sharesmb="rw,ro=sf.example.com"`
	`resource_name`	The share name is the resource name and is available for reading and writing for all clients if there are no SMB exceptions.	`sharesmb=myshare`
	`resource_name``,rw`	The share name is the resource name and is available for reading and writing for all clients except those for which the `ro` exception is defined. SMB exceptions may or may not be defined.	`sharesmb="myshare,rw"` or `sharesmb="myshare,rw,ro=sf.example.com"`
Read only	`ro`	The share name is the dataset name and is available for reading only for all hosts except those for which the `rw` exception is defined.	`sharesmb="ro,rw=sf.example.com"`
	`resource_name``,ro`	The share name is the resource name and is available for reading only for all clients except those for which the `rw` exception is defined. SMB exceptions may or may not be defined.	`sharesmb="myshare,ro"` or `sharesmb="myshare,ro,rw=sf.example.com"`

The following example sets the share mode for all clients to read-only.

```
set sharesmb=ro
```

Additional SMB exceptions can be specified by appending text of the form `option=collection`, where *option* is either `ro` or `rw`. You cannot grant root access with SMB exceptions. The collection is specified by the prefix character from table 114, and either a DNS hostname/domain name or CIDR network number.

For example, to grant read-write access to all hosts in the `sf.example.com` domain:

```
set sharesmb="ro,rw=.sf.example.com"
```

This example grants read-only access to clients with the IP addresses 2001:db8:410:d43::/64 and 192.0.2.254/22:

```
set sharesmb="on,ro=@[2001:db8:410:d43::/64]:@192.0.2.254/22"
```

Netgroup names can be used anywhere an individual fully qualified hostname can be used. For example, you can permit read-write access to the `engineering` netgroup as follows:

```
set sharesmb="ro,rw=engineering"
```

Share-Level ACLs

A share-level access control list (ACL), when combined with the ACL of a file or directory in the share, determines the effective permissions for that file. By default, this ACL grants everyone full control. This ACL provides another layer of access control above the ACLs on files and allows for more sophisticated access control configurations. This property may only be set once the filesystem has been exported by configuring the SMB resource name. If the filesystem is not exported over the SMB protocol, setting the share-level ACL has no effect.

When access-based enumeration is enabled, clients may see directory entries for files which they cannot open. Directory entries are filtered only when the client has no access to that file. For example, if a client attempts to open a file for read/write access but the ACL grants only read access, that open request will fail but that file will still be included in the list of entries.

For more information about ACLs, see Access Control Lists for Filesystems.

HTTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the HTTP protocol (`sharedav`) and for an Object Store, users can set the share mode to determine if the filesystem is available for read only (`ro`), read and write (`rw` or `on`), or neither (`off`).

Related Topics

- Project Properties
- Filesystem Properties

FTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the FTP protocol (`shareftp`), users can set the share mode to determine if the filesystem is available for read only (`ro`), read and write (`rw` or `on`), or neither (`off`).

Related Topics

- Project Properties
- Filesystem Properties

SFTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the SFTP protocol (`sharesftp`), users can set the share mode to determine if the filesystem is available for read only (`ro`), read and write (`rw` or `on`), or neither (`off`).

Related Topics

- Project Properties
- Filesystem Properties

TFTP Protocol

Each share has protocol-specific properties that define the behavior of different protocols for that share. These properties can be defined for each share or inherited from a share's project. For the TFTP protocol (`sharetftp`), users can set the share mode to determine if the filesystem is available for read only (`ro`), read and write (`rw` or `on`), or neither (`off`).

Related Topics

- Project Properties
- Filesystem Properties

Access Control Lists for Filesystems

You can set options to control ACL behavior as well as control access to the root directory of the filesystem.

> **Note:**
>
> ACLs are available only for filesystems.

For more information about ACLs, see the following topics:

- Root Directory Access
- ACL Behavior on Mode Change
- ACL Inheritance Behavior
- Root Directory ACL

Root Directory Access

To set basic access control for the root directory of the filesystem, go to **Shares: Shares: *filesystem*: Access**. These settings can be managed in-band via whichever protocols are being used, but they can also be specified here for convenience. These properties cannot be changed on a read-only filesystem, as they require changing metadata for the root directory of the filesystem.

- **User** - The owner of the root directory. This can be specified as a user ID or user name. For more information on mapping UNIX and Windows users, see Identity Mapping. For UNIX-based NFS access, this can be changed from the client using the `chown` command.
- **Group** - The group of the root directory. This can be specified as a group ID or group name. For more information on mapping UNIX and Windows groups, see Identity Mapping. For UNIX-based NFS access, this can be changed from the client using the `chgrp` command.
- **Permissions** - Standard UNIX permissions for the root directory. For UNIX-based NFS access, this can be changed from the client using the `chmod` command. The permissions are divided into three types.

Access Type	Description
User	User that is the current owner of the directory.
Group	Group that is the current group of the directory.
Other	All other accesses.

For each access type, the following permissions can be granted.

Type	Description
Read (R)	Permission to list the contents of the directory.
Write (W)	Permission to create files in the directory.
Execute (X)	Permission to look up entries in the directory. If users have execute permissions but not read permissions, they can access files explicitly by name but not list the contents of the directory.

Related Topics

- ACL Behavior on Mode Change
- ACL Inheritance Behavior
- Root Directory ACL

ACL Behavior on Mode Change

When an ACL is modified via `chmod`(2) using the standard UNIX user/group/other permissions, the simplified mode change request will interact with the existing ACL in different ways depending on the setting of this property. To edit the ACL behavior on mode change, see Editing a Project - BUI, CLI.

Table 4-61 Mode Change Values

BUI Value	CLI Value	Description
Discard ACL	`discard`	All ACL entries that do not represent the mode of the directory or file are discarded. This is the default behavior.
Mask ACL with mode	`mask`	The permissions are reduced, such that they are no greater than the group permission bits, unless it is a user entry that has the same UID as the owner of the file or directory. In this case, the ACL permissions are reduced so that they are no greater than owner permission bits. The mask value also preserves the ACL across mode changes, provided an explicit ACL set operation has not been performed.

Table 4-61 (Cont.) Mode Change Values

BUI Value	CLI Value	Description
Do not change ACL	`passthrough`	No changes are made to the ACL other than generating the necessary ACL entries to represent the new mode of the file or directory.

Related Topics

- Root Directory ACL
- ACL Inheritance Behavior
- Root Directory ACL

ACL Inheritance Behavior

When a new file or directory is created, it is possible to inherit existing ACL settings from the parent directory. This property controls how this inheritance works. These property settings usually only affect ACL entries that are flagged as inheritable - other entries are not propagated regardless of this property setting. However, all trivial ACL entries are inheritable when used with SMB. A trivial ACL represents the traditional UNIX `owner/group/othe`r entries. To edit the ACL inheritance behavior, see Editing a Project - BUI, CLI.

Table 4-62 ACL Inheritance Behavior Values

BUI Value	CLI Value	Description
Do not inherit entries	`discard`	No ACL entries are inherited. The file or directory is created according to the client and protocol being used.
Only inherit deny entries	`noallow`	Only inheritable ACL entries specifying `deny` permissions are inherited.
Inherit all but "write ACL" and "change owner"	`restricted`	Removes the `write_acl` and `write_owner` permissions when the ACL entry is inherited, but otherwise leaves inheritable ACL entries untouched. This is the default.
Inherit all entries	`passthrough`	All inheritable ACL entries are inherited. The `passthrough` mode is typically used to cause all `data` files to be created with an identical mode in a directory tree. An administrator sets up ACL inheritance so that all files are created with a mode, such as 0664 or 0666.
Inherit all but "execute" when not specified	`passthrough-x`	Same as `passthrough` except that the owner, group, and everyone ACL entries inherit the execute permission only if the file creation mode also requests the execute bit. The `passthrough` setting works as expected for data files, but you might want to optionally include the execute bit from the file creation mode into the inherited ACL. One example is an output file that is generated from tools, such as `cc` or `gcc`. If the inherited ACL does not include the execute bit, then the output executable from the compiler won't be executable until you use `chmod(1)` to change the file's permissions.

Table 4-62 (Cont.) ACL Inheritance Behavior Values

BUI Value	CLI Value	Description
Inherit all, but preserve mode from client	`passthrough-mode-preserve`	Inheritable ACL entries are inherited, while preserving the creation mode specified by the application. This preserves the inheritance bits so SMB creates ACLs that interoperate well with shares accessed over NFS and SMB simultaneously. This property setting is only available after applying the deferred update for ACL Passthrough with Mode Preservation. For more information, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

When using SMB to create a file in a directory with a trivial ACL, all ACL entries are inherited. As a result, the following behavior occurs:

- Inheritance bits display differently when viewed in SMB or NFS. When viewing the ACL directory in SMB, inheritance bits are displayed. In NFS, inheritance bits are not displayed.
- When a file is created in a directory using SMB, its ACL entries are shown as inherited; however, when viewed through NFS, the directory has no inheritable ACL entries.
- If the ACL is changed so that it is no longer trivial, for example, by adding an access control entry (ACE), this behavior does not occur.
- If the ACL is modified using SMB, the resulting ACL will have the previously synthetic inheritance bits turned into real inheritance bits.

Related Topics

Project Properties

Root Directory ACL

Fine-grained access on files and directories is managed via Access Control Lists. An ACL describes which permissions are granted, if any, to specific users or groups. Oracle ZFS Storage Appliance supports NFSv4.0 and NFSv4.1-style ACLs, also accessible over SMB. POSIX draft ACLs (used by NFSv3) are not supported. Some trivial ACLs can be represented over NFSv3, but making complicated ACL changes may result in undefined behavior when accessed over NFSv3.

Like root directory access, this property only affects the root directory of the filesystem. ACLs can be controlled through in-band protocol management; BUI and CLI provide a way to set the ACL just for the root directory of the filesystem. You can use in-band management tools if the BUI is not an option. Changing this ACL does not affect existing files and directories in the filesystem. Depending on the ACL inheritance behavior, these settings may or may not be inherited by newly created files and directories. However, all ACL entries are inherited when SMB is used to create a file in a directory with a trivial ACL.

An ACL is composed of any number of ACEs (access control entries). Each ACE describes a type/target, a set of permissions, inheritance flags and a type. ACEs are applied in order, starting at the beginning of the ACL, to determine whether a given action should be permitted. For information on in-band configuration ACLs through data protocols, consult the appropriate client documentation. The BUI interface for managing ACLs and the effect on the root directory are described here.

Table 4-63 Share - ACL Types

Type	Description
Owner	Current owner of the directory. If the owner is changed, this ACE will apply to the new owner.
Group	Current group of the directory. If the group is changed, this ACE will apply to the new group.
Everyone	Any user.
Named User	User named by the `target` field. The user can be specified as a user ID or a name resolvable by the current name service configuration.
Named Group	Group named by the `target` field. The group can be specified as a group ID or a name resolvable by the current name service configuration.

Table 4-64 Share - ACE Types

ACE Type	Description
Allow	The permissions are explicitly granted to the ACE target.
Deny	The permissions are explicitly denied to the ACE target.
Audit	Generates an audit record based on the permissions field. As with the Allow and Deny ACE types, the action is a success (`S`) or a failure (`F`). Because this ACL is only for the filesystem root directory, the inheritance is usually set to file and directory. Audit records are sent via the Syslog Relay service. To send audit events to a remote target, set audit class option Per File Audit (`PerFileAudit`). See Syslog Configuration.

Table 4-65 Share - ACL Permissions

Permission	Description
(`r`) Read Data/List Directory	Permission to list the contents of a directory. When inherited by a file, permission to read the data of the file.
(`x`) Execute File/Traverse Directory	Permission to traverse (lookup) entries in a directory. When inherited by a file, permission to execute the file.
(`a`) Read Attributes	Permission to read basic attributes (non-ACLs) of a file. Basic attributes are considered to be the stat level attributes, and allowing this permission means that the user can execute `ls` and `stat` equivalents.
(`R`) Read Extended Attributes	Permission to read the extended attributes of a file or do a lookup in the extended attributes directory.
(`w`) Write Data/Add File	Permission to add a new file to a directory. When inherited by a file, permission to modify a file's data anywhere in the file's offset range. This include the ability to grow the file or write to any arbitrary offset.
(`p`) Append Data/Add Subdirectory	Permission to create a subdirectory within a directory. When inherited by a file, permission to modify the file's data, but only starting at the end of the file. This permission (when applied to files) is not currently supported.
(`d`) Delete	Permission to delete a file.

Table 4-65 (Cont.) Share - ACL Permissions

Permission	Description
(D) Delete Child	Permission to delete a file within a directory. As of the 2011.1 software release, if the sticky bit is set, a child file can only be deleted by the file owner.
(A) Write Attributes	Permission to change the times associated with a file or directory.
(W) Write Extended Attributes	Permission to create extended attributes or write to the extended attributes directory.
(c) Read ACL/Permissions	Permission to read the ACL.
(C) Write ACL/Permissions	Permission to write the ACL or change the basic access types.
(o) Change Owner	Permission to change the owner.
(f) Apply to Files	Inherit to all newly created files in a directory.
(d) Apply to Directories	Inherit to all newly created directories in a directory.
(i) Do not apply to self	The current ACE is not applied to the current directory, but does apply to children. This flag requires one of `Apply to Files` or `Apply to Directories` to be set.
(n) Do not apply past children	The current ACE should only be inherited one level of the tree, to immediate children. This flag requires one of `Apply to Files` or `Apply to Directories` to be set.

When the option to use Windows default permissions is used at share creation time, an ACL with the following three entries is created for the share's root directory:

Table 4-66 Share Root Directory Entities

Type	Action	Access
Owner	Allow	Full Control
Group	Allow	Read and Execute
Everyone	Allow	Read and Execute

In the CLI, set the root directory ACL properties after navigating to the `shares` context and selecting a project and filesystem. Use colons to separate the ACE properties, and separate multiple ACE entries with commas. The `target` and `inheritance` fields are optional. To set the properties, enter `set root_acl=ace1,ace2,ace3,...`, where *`acen`* is:

```
type:[target:]permissions:[inheritance:]type
```

Examples:

- `set root_acl=owner@:r:allow`
- `set root_acl=everyone@:rwx:fd:allow`
- `set root_acl=user:root:r:allow`

Working with Schemas

In addition to the standard built-in properties, you can configure any number of additional properties that are available on all shares and projects. These properties are given basic types for validation purposes, and are inherited like most other standard properties. The values are never consumed by the software in any way, and exist solely for end-user consumption. The property schema is global to the system, across all pools, and is synchronized between cluster peers.

To work with schemas, see the following sections:

- Creating a Schema (BUI)
- Creating a Schema (CLI)
- Schema Properties

Creating a Schema (BUI)

Use the following procedure create a schema.

1. From the **Shares** menu, select **Schema**.
2. Click the add icon ⊕ next to **Properties** to add a new property to the schema property list.
3. Enter the **Name** of the property (contact).
4. Enter a **Description** of the property ("Owner Contact").
5. Choose a **Type** for the new property (**Email Address**).
6. Click **APPLY**.
7. Navigate to an existing share or project.
8. Change the **Owner Contact** property under the **Custom Properties** section.

Creating a Schema (CLI)

Use the following procedure create a schema.

1. Go to `shares schema`.
2. To create a new property named contact, enter `create contact`.
3. Set the description for the property.

   ```
   hostname:shares schema> set description="Owner Contact"
   ```

4. Set the type of the property.

   ```
   hostname:shares schema> set type=EmailAddress
   ```

5. Enter `commit`.
6. Go to an existing share or project.
7. Set the custom contact property, and enter `commit`.

   ```
   hostname:shares default> set custom:owner=chris@corp
                      custom:owner = chris@corp (uncommitted)
   hostname:shares default> commit
   ```

Example 4-1 Example Schema

The schema context can be found at `shares schema`.

```
hostname:> shares schema
hostname:shares schema> show
Properties:

NAME          TYPE          DESCRIPTION
owner         EmailAddress  Owner Contact
```

Each property is a child of the schema context, using the name of the property as the token. To create a property, use the `create` command:

```
hostname:shares schema> create department
hostname:shares schema department (uncommitted)> get
                          type = String
                   description = department
hostname:shares schema department (uncommitted)> set description="Department Code"
                   description = Department Code (uncommitted)
hostname:shares schema department (uncommitted)> commit
hostname:shares schema>
```

Within the context of a particular property, fields can be set using the standard CLI commands:

```
hostname:shares schema> select owner
hostname:shares schema owner> get
                          type = EmailAddress
                   description = Owner Contact
hostname:shares schema owner> set description="Owner Contact Email"
                   description = Owner Contact Email (uncommitted)
hostname:shares schema owner> commit
```

After custom properties have been defined, they can be accessed like any other property under the name `custom:`*`property_name`*:

```
hostname:shares default> get
...
             custom:department = 123-45-6789
                  custom:owner =
...
hostname:shares default> set custom:owner=chris@corp
                  custom:owner = chris@corp (uncommitted)
hostname:shares default> commit
```

Schema Properties

To define custom properties, navigate to **Shares: Schema** in the BUI or `shares schema` in the CLI. The current schema is displayed as a list, and entries can be added or removed as needed. Each property has the following fields:

Table 4-67 Schema Property Fields

Field	Description
NAME	The CLI name for this property. This must contain only alphanumeric characters or the characters `.:_\`
DESCRIPTION	The BUI name for this property. This can contain arbitrary characters and is used in the help section of the CLI.

Table 4-67   (Cont.) Schema Property Fields

Field	Description
`TYPE`	The property type, for validation purposes. This must be one of the types described in the following table.

The valid types for properties are:

Table 4-68   Valid Types for Properties

BUI Type	CLI Type	Description
String	`String`	Arbitrary string data. This is the equivalent of no validation.
Integer	`Integer`	A positive or negative integer.
Positive Integer	`PositiveInteger`	A positive integer.
Boolean	`Boolean`	A `true` or `false` value. In the BUI, this is presented as a check box, while in the CLI, it must be one of the values `true` or `false`.
Email Address	`EmailAddress`	An email address. Only minimal syntactic validation is done.
Hostname or IP	`Host`	A valid DNS hostname or IP (v4 or v6) address.

Once defined, the properties are available under the **General** properties tab, using the description provided in the property table. Properties are identified by their CLI name, so renaming a property will have the effect of removing all existing settings on the system. A property that is removed and later renamed back to the original name will still refer to the previously set values. Changing the types of properties, while supported, may have undefined results on existing properties on the system. Existing properties will retain their current settings, even if they would be invalid given the new property type.

5

Shadow Migration

A common task for administrators is to move data from one location to another. In the most abstract sense, this problem encompasses a large number of use cases, from replicating data between servers to keeping user data on laptops in synchronization with servers. There are many external tools available to do this, but Oracle ZFS Storage Appliance has two integrated solutions for migrating data that addresses the most common use cases. The first, replication, is intended for replicating data between one or more appliances, and is covered separately; see Remote Replication. The second, shadow migration, is described in this section.

Shadow migration is a process for migrating data from external NAS sources with the intent of replacing or decommissioning the original once the migration is complete. This is most often used when introducing a new appliance into an existing environment in order to take over file sharing duties of another server, but a number of other novel uses are possible.

To use shadow migration, see the following sections:

- Understanding Shadow Migration
- Creating a Shadow Filesystem
- Managing Background Migration
- Handling Migration Errors
- Monitoring Migration Progress
- Canceling Migration
- Snapshotting Shadow File Systems
- Backing Up Shadow File Systems
- Replicating Shadow File Systems
- Migrating Local File Systems
- Using Shadow Migration Analytics
- Testing Potential Shadow Migration using the CLI
- Migrating Data from an Active NFS Server using the CLI

Understanding Shadow Migration

Shadow migration uses interposition, but it is integrated into Oracle ZFS Storage Appliance and does not require a separate physical machine. When filesystems are created, they can optionally "shadow" an existing directory, either locally or over NFS. In this scenario, downtime is scheduled once, where the source appliance X is placed into read-only mode, a share is created with the shadow property set, and clients are updated to point to the new share on the appliance. Clients can then access the appliance in read-write mode.

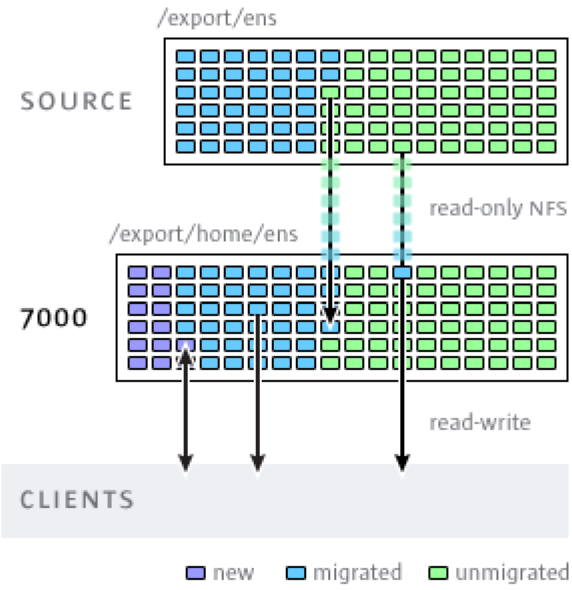

After the shadow property is set, data is transparently migrated in the background from the source appliance locally. If a request comes from a client for a file that has not yet been migrated, the appliance will automatically migrate this file to the local server before responding to the request. This may incur some initial latency for some client requests, but once a file has been migrated, all accesses are local to the appliance and have native performance. It is often the case that the current working set for a filesystem is much smaller than the total size, so once this working set has been migrated, regardless of the total native size on the source, there will be no perceived impact on performance.

The downside to shadow migration is that it requires a commitment before the data has finished migrating, though this is the case with any interposition method. During the migration, portions of the data exist in two locations, which means that backups are more complicated, and snapshots may be incomplete and/or exist only on one host. It is therefore extremely important that any migration between two hosts first be tested thoroughly to make sure that identity management and access controls are setup correctly. This need not test the entire data migration, but it should be verified that files or directories that are not world readable are migrated correctly, ACLs (if any) are preserved, and identities are properly represented on the new system.

Shadow migration is implemented using on-disk data within the filesystem, so there is no external database and no data stored locally outside the storage pool. If a pool is failed over in a cluster, or both system disks fail and a new head node is required, all data necessary to continue shadow migration without interruption will be kept with the storage pool.

The following lists the restrictions on the shadow source:

- To properly migrate data, the source filesystem or directory **must be read-only**. Changes made to files source may or may not be propagated based on timing, and changes to the directory structure can result in unrecoverable errors on the appliance.
- Shadow migration supports migration only from NFS sources. NFSv4.0 and NFSv4.1 filesystems will yield the best results. NFSv2 and NFSv3 migration are possible, but ACLs will be lost in the process and files that are too large for NFSv2 cannot be migrated using that protocol. Migration from SMB sources is not supported.
- Shadow migration of LUNs is not supported.

During migration, if the client accesses a file or directory that has not yet been migrated, there is an observable effect on behavior. The following lists the shadow file system semantics:

- For directories, client requests are blocked until meta-data infrastructure is created on the migration target for any intervening directories for which infrastructure is not yet established. For files, only the portion of the file being requested is migrated, and multiple clients can migrate different portions of a file at the same time.
- Files and directories can be arbitrarily renamed, removed, or overwritten on the shadow filesystem without any effect on the migration process.
- For files that are hard links, the hard link count may not match the source until the migration is complete.
- The majority of file attributes are migrated when the directory is created, but the on-disk size (`st_nblocks` in the UNIX `stat` structure) is not available until a read or write operation is done on the file. The logical size will be correct, but a `du`(1) or other command will report a zero size until the file contents are actually migrated.
- If the appliance is rebooted, the migration will pick up where it left off originally. While it will not have to re-migrate data, it may have to traverse some already-migrated portions of the local filesystem, so there may be some impact to the total migration time due to the interruption.
- Data migration makes use of private extended attributes on files. These are generally not observable except on the root directory of the filesystem or through snapshots. Adding, modifying, or removing any extended attribute that begins with `SUNWshadow` will have undefined effects on the migration process and will result in incomplete or corrupt state. In addition, filesystem-wide state is stored in the `.SUNWshadow` directory at the root of the filesystem. Any modification to this content will have a similar effect.
- After a filesystem has completed migration, an alert will be posted, and the shadow attribute will be removed, along with any applicable metadata. After this point, the filesystem will be indistinguishable from a normal filesystem.
- Data can be migrated across multiple filesystems into a singe filesystem, through the use of NFSv4.0 or NFSv4.1 automatic client mounts (sometimes called "mirror mounts") or nested local mounts.

Use the following rules to migrate identity information for files, including ACLs:

- The migration source and target appliance must have the same name service configuration.
- The migration source and target appliance must have the same NFSv4.0 or NFSv4.1 `mapid` domain
- The migration source must support NFSv4.0 and NFSv4.1. Use of NFSv3 is possible, but some loss of information will result. Basic identity information (owner and group) and POSIX permissions will be preserved, but any ACLs will be lost.
- The migration source must be exported with root permissions to the appliance.

If you see files or directories owned by `nobody`, it likely means that the appliance does not have name services setup correctly, or that the NFSv4.0 or NFSv4.1 `mapid` domain is different. If you get `permission denied` errors while traversing filesystems that the client should otherwise have access to, the most likely problem is failure to export the migration source with root permissions.

Creating a Shadow Filesystem

The shadow migration source can only be set when a filesystem is created. In the BUI, this is available in the filesystem creation dialog box. In the CLI, it is available as the `shadow` property. The property takes one of the following forms:

- **Local** - `file:///<path>`
- **NFS** - `nfs://<host>/<path>`

The BUI also allows the alternate form `<host>:/<path>` for NFS mounts, which matches the syntax used in UNIX systems. The BUI also sets the protocol portion of the setting (`file://` or `nfs://`) via the use of a pull-down menu. When creating a filesystem, the server will verify that the path exists and can be mounted.

Managing Background Migration

When a share is created, it will automatically begin migrating in the background, in addition to servicing inline requests. This migration is controlled by the shadow migration service. There is a single global tunable, which is the number of threads dedicated to this task. Increasing the number of threads will result in greater parallelism at the expense of additional resources.

The shadow migration service can be disabled, but this should only be used for testing purposes, or when the active of shadow migration is overwhelming the system to the point where it needs to be temporarily stopped. When the shadow migration service is disabled, synchronous requests are still migrated as needed, but no background migration occurs. With the service disabled, no shadow migration will ever complete, even if all the contents of the filesystem are read manually. It is highly recommended to always leave the service enabled.

Handling Migration Errors

Because shadow migration requires committing new writes to the server prior to migration being complete, it is very important to test migration and monitor for any errors. Errors encountered during background migration are kept and displayed in the BUI as part of shadow migration status. Errors encountered during other synchronous migration are not tracked, but will be accounted for once the background process accesses the affected file. For each file, the remote filename as well as the specific error are kept. Clicking on the information icon next to the error count will bring up this detailed list. The error list is not updated as errors are fixed, but simply cleared by virtue of the migration completing successfully.

Shadow migration will not complete until all files are migrated successfully. If there are errors, the background migration will continually retry the migration until it succeeds. This allows the administrator to fix any errors (such as permission problems), let the migration complete, and be assured of success. If the migration cannot complete due to persistent errors, the migration can be canceled, leaving the local filesystem with whatever data was able to be migrated. This should only be used as a last resort; once migration has been canceled, it cannot be resumed.

Monitoring Migration Progress

To monitor shadow migration progress, Oracle ZFS Storage Appliance provides such statistics as:

- Size of data transferred so far
- Estimate of remaining size to be migrated

- Migration time so far
- Migration time remaining
- Migration errors

At the beginning of migration, the appliance obtains the source filesystem statistics and calculates its size. It uses these values to provide a reasonably accurate visual representation of migration progress and an estimation of the remaining data to be migrated. Of note, the remaining bytes is an estimate based on the assumption that an entire filesystem is being migrated. If only part of the source file system is migrated, the remaining bytes estimate is inaccurate. If the source filesystem has nested filesystems, the total filesystem size is recalculated when the nested mount point is discovered during migration, and the remaining bytes are re-estimated based on the newly calculated total. Estimation of remaining bytes may be inaccurate if the source filesystem uses compression. These values are available in the BUI and CLI through both the standard filesystem properties as well as properties of the shadow migration node (or UI panel).

> **Note:**
>
> When a sparse file (a file with empty blocks) is migrated, the target file will be smaller than the source file size. Shadow migration does not write the empty blocks to the target file, resulting in less space usage.

The following tasks describe how to monitor shadow migration progress and view any resulting errors. To view shadow migration errors using the RESTful API, see Filesystem Operations in *Oracle ZFS Storage Appliance RESTful API Guide, Release OS8.8.x*.

Monitoring Migration Progress and Errors (BUI)

Use the following procedure to monitor migration progress and errors.

1. From the **Shares** menu, select **Shares**, and select a filesystem with shadow migration source.
2. Under **Shadow Migration**, examine the progress bar and shadow migration status.

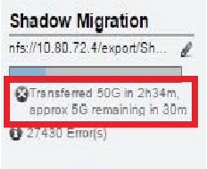

Shadow Migration

nfs://10.80.72.4/export/Sh...

Transferred 50G in 2h34m, approx 5G remaining in 30m

27430 Error(s)

3. To view shadow migration errors, click the edit icon ✎ .

Monitoring Migration Progress and Errors (CLI)

Use the following procedure to monitor migration progress and errors.

1. Go to a filesystem with shadow migration source, and enter `shadow`, and then enter `list`.

```
hostname:shares default/file_sys1> shadow
hostname:shares default/file_sys1 shadow> list
Properties:
                              source = nfs://zfs0000-15/sm/errors
```

```
                   transferred = (unset)
                     remaining = 1.37G
                       elapsed = 0h3m
                        errors = 23
                      complete = true
                          time = (unset)

Children:
                           errors => Shadow Migration Errors
hostname:shares default/file_sys1 shadow>
```

> **Note:**
>
> If there is no shadow migration source defined (`shadow=none`), then the `shadow` command is invalid for the filesystem:
>
> ```
> hostname:shares default/xyz_1> shadow
> error: invalid command "shadow"
> ```

2. To view shadow migration errors, enter child node `errors`, and enter `help` to view a list of subcommands.

```
hostname:shares default/file_sys1 shadow> errors

hostname:shares default/file_sys1 shadow errors> help

Subcommands that are valid in this context:

   help [topic]         => Get context-sensitive help. If [topic] is specified,
                           it must be one of "builtins", "commands", "general",
                           "help" or "script".

   show                 => Show information pertinent to the current context

   done                 => Finish operating on "errors"

   select [entry]       => Select the specified entry to get its properties,
                           set its properties, or run a subcommand

   list                 => Lists up to the first 100 errors. The "-a" option may be
                           used to list all the errors if there are more that 100
                           errors. The "-number" option may be used to list the first
                           (number) of errors. Format: list -a or list -xx
```

3. Enter `show` to view individual migration errors in the current context.

```
hostname:shares default/file_sys1 shadow errors> show
Errors:

PATH                             REASON
ak-2013-dev-on-clone.tar.gz     Permission denied
test_dir/CREDITS.html           Permission denied
test_dir/CREDITS_ja.html        Permission denied
test_dir/CREDITS_pt_BR.html     Permission denied
test_dir/CREDITS_zh_CN.html     Permission denied
test_dir/DISTRIBUTION.txt       Permission denied
test_dir/LEGALNOTICE.txt        Permission denied
test_dir/LICENSE.txt            Permission denied
test_dir/README.html            Permission denied
test_dir/README_ja.html         Permission denied
```

```
test_dir/README_pt_BR.html      Permission denied
test_dir/README_zh_CN.html      Permission denied
test_dir/THIRDPARTYLICENSE.txt  Permission denied
test_dir/bin                    Permission denied
test_dir/cnd2                   Permission denied
test_dir/etc                    Permission denied
test_dir/gsf1                   Permission denied
test_dir/ide10                  Permission denied
test_dir/modCluster.properties  Permission denied
test_dir/nb6.5                  Permission denied
test_dir/forms.css              Permission denied
test_dir/platform9              Permission denied
test_dir/websvccommon1          Permission denied
```

4. To view an individual error, enter the `select` command and an individual error name. Then enter `show`.

   To view individual error properties, use the `get` command.

   ```
   hostname:shares default/file_sys1 shadow errors> select test_dir/nb6.5

   hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> show
   Properties:
                             path = test_dir/nb6.5
                           reason = Permission denied

   hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> get path
                             path = test_dir/nb6.5
   hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> get reason
                           reason = Permission denied
   hostname:shares default/file_sys1 shadow errors test_dir/nb6.5> done
   hostname:shares default/file_sys1 shadow errors>
   ```

Canceling Migration

Migration can be canceled, but this should only be done in extreme circumstances when the source is no longer available. Once migration has been canceled, it cannot be resumed. The primary purpose is to allow migration to complete when there are uncorrectable errors on the source. If the complete filesystem has finished migration except for a few files or directories, and there is no way to correct these errors (that is, the source is permanently broken), then canceling the migration will allow the local filesystem to resume status as a 'normal' filesystem.

To cancel migration in the BUI, click the close icon ⊗ next to the progress bar in the left column of the share in question. In the CLI, navigate to the `shadow` node beneath the filesystem, and run the `cancel` command.

Snapshotting Shadow File Systems

Shadow filesystems can be snapshotted; however, the state of what is included in the snapshot is arbitrary. Files that have not yet been migrated will not be present, and implementation details (such as `SUNWshadow` extended attributes) may be visible in the snapshot. This snapshot can be used to restore individual files that have been migrated or modified since the original migration began. Because of this, it is recommended that any snapshots be kept on the source until the migration is completed, so that file that have not been migrated can still be retrieved from the source if necessary. Depending on the retention policy, it may be necessary to extend retention on the source in order to meet service requirements. While snapshots can be taken, these snapshots cannot be rolled back to, nor can they be the source of a clone.

Backing Up Shadow File Systems

Filesystems that are actively migrating shadow data can be backed using NDMP, as with any other filesystem. The shadow setting is preserved with the backup stream, but will be restored only if a complete restore of the filesystem is done and the share does not already exist. Restoring individual files from such a backup stream or restoring into existing filesystems may result in inconsistent state or data corruption. During the full filesystem restore, the filesystem will be in an inconsistent state (beyond the normal inconsistency of a partial restore) and shadow migration will not be active. Only when the restore is completed is the shadow setting restored. If the shadow source is no longer present or has moved, the administrator can observe any errors and correct them as necessary.

Replicating Shadow File Systems

Filesystems that are actively migrating shadow data can be replicated using the normal mechanism, but only the migrated data is sent in the data stream. As such, the remote side contains only partial data that may represent an inconsistent state. The shadow setting is sent along with the replication stream, so when the remote target is failed over, it will keep the same shadow setting. As with restoring an NDMP backup stream, this setting may be incorrect in the context of the remote target. After failing over the target, the administrator can observe any errors and correct the shadow setting as necessary for the new environment.

Migrating Local Filesystems

In addition to its primary purpose of migrating data from remote sources, the same mechanism can also be used to migrate data from a local filesystem to another local filesystem on Oracle ZFS Storage Appliance. This can be used to change settings that otherwise cannot be modified, such as creating a compressed version of a filesystem, or changing the recordsize for a filesystem after the fact. In this model, the old share (or subdirectory within a share) is made read-only or moved aside, and a new share is created with the shadow property set using the `file` protocol. Clients access this new share, and data is written using the settings of the new share.

Using Shadow Migration Analytics

In addition to standard monitoring on a per-share basis, it is also possible to monitor shadow migration system-wide using analytics. The shadow migration analytics are available under the `Data Movement` category. There are three basic statistics available:

- **Shadow Migration Requests -** This statistic tracks requests for files or directories that are not cached and known to be local to the filesystem. It does account for both migrated and unmigrated files and directories, and it can be used to track the latency incurred as part of shadow migration, as well as track the progress of background migration. It can be broken down by file, share, project, or latency. It currently encompasses both synchronous and asynchronous (background) migration, so it is not possible to view only latency visible to clients.
- **Shadow Migration Bytes -** This statistic tracks bytes transferred as part of migrating file or directory contents. This does not apply to metadata (extended attributes, ACLs, and so on). It gives a rough approximation of the data transferred, but source datasets with a large amount of metadata will show a disproportionally small bandwidth. The complete bandwidth can be observed by looking at network analytics. This statistic can be broken down by local filename, share, or project.

- **Shadow Migration Operations** - This statistic tracks operations that require going to the source filesystem. This can be used to track the latency of requests from the shadow migration source. It can be broken down by file, share, project, or latency.

Testing Potential Shadow Migration using the CLI

Before attempting a complete migration, it is important to test the migration to make sure that Oracle ZFS Storage Appliance has appropriate permissions and security attributes are translated correctly. After you are confident that the basic setup is functional, the filesystems can be set up for the final migration.

> **Note:**
>
> As part of capacity planning, remember to take into account default/user group quotas because the quotas could be exceeded if the source is larger than the destination. Also, shadow migration will fail if the target runs out of disk space.

1. Configure the source so that the appliance has root access to the share. This typically involves adding an NFS host-based exception or setting the anonymous user mapping (the latter having more significant security implications).
2. Create a share on the local filesystem with the shadow attribute set to `nfs://<host>/<snapshotpath>` in the CLI, or `<host>/<snapshotpath>` in the BUI (with the `NFS` protocol selected). The snapshot should be a read-only copy of the source. If no snapshots are available, a read-write source can be used, but may result in undefined errors.
3. Validate that file contents and identity mapping are correctly preserved by traversing the file structure.
4. If the data source is read-only (as with a snapshot), let the migration complete, and verify that there were no errors in the transfer.

Migrating Data from an Active NFS Server using the CLI

Use the following procedure to migrate data from an active NFS server using the CLI. Note that shadow migration fails if it encounters files under `procfs` or the following special file types: doors, sockets, and event ports.

1. Schedule downtime during which clients can be quiesced and reconfigured to point to a new server.
2. Configure the source so that Oracle ZFS Storage Appliance has root access to the share. This typically involves adding an NFS host-based exception, or setting the anonymous user mapping (the latter having more significant security implications).
3. Configure the source to be read-only. This step is technically optional, but it is much easier to guarantee compliance if it is impossible for misconfigured clients to write to the source while migration is in progress.
4. Create a share on the local filesystem with the shadow attribute set to `nfs://<host>/<path>` in the CLI, or `<host>/<path>` in the BUI (with the `NFS` protocol selected).
5. Reconfigure clients to point at the local share on the appliance.

   At this point, shadow migration should be running in the background, and client requests should be serviced as necessary. You can observe the progress as described earlier.

Multiple filesystems can be created during a single scheduled downtime through scripting the CLI.

6 Snapshots and Clones

> **Note:**
>
> Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Using snapshots and clones, you can make point-in-time copies of a share or project. These copies can be useful as backups or as different working versions.

A snapshot is a read-only copy of a filesystem, LUN, or project. Taking a project snapshot is equivalent to snapshotting all of the shares in the project. A snapshot takes up no additional space when it is first created, but as the active share changes, the snapshot takes up additional space, with a maximum equivalent to the size of the share at the time the snapshot was taken.

A clone is a writable copy of a filesystem or LUN snapshot and can be treated as an independent share. Clones of projects are not supported. Like a snapshot, a clone consumes no additional space when it is first created, but as new data is written to the clone, the space required for the changes are associated with the clone.

You can take snapshots manually, or you can set a schedule so that snapshots are taken automatically every half-hour, hour, day, week, or month. Some snapshots are taken by the appliance automatically during replication updates; these will appear on the snapshot page with `.ndmp` and `.rr` in their names.

> **Note:**
>
> When the file retention feature, which is different from the snapshot retention feature, is used with snapshot rollback, certain restrictions can apply. Rollback can be performed on a filesystem with the privileged file retention policy, even when unexpired retained files exist. Filesystems with the mandatory file retention policy can never be rolled back, even when all retained files have expired. For more information, see File Retention Management.

For information about snapshot space management, see:Snapshot Space Management.

To take manual snapshots or schedule automatic snapshots of projects or shares, use the following tasks:

- Taking a Snapshot - BUI, CLI
- Scheduling Snapshots - BUI, CLI
- Setting a Scheduled Snapshot Label - BUI, CLI

You can make clones of a snapshot, which can be useful to create numerous working versions of one share. To make clones, use the following tasks:

- Cloning a Snapshot - BUI, CLI
- Cloning a Clone
- Cloning a Replication Package - BUI, CLI

To determine the relationships between existing snapshots and clones, use the following tasks:

- Viewing Clones of a Snapshot - BUI, CLI
- Viewing a Clone Origin - BUI, CLI

To view and edit existing snapshots, snapshot schedules, and snapshot retention policies, use the following tasks:

- Viewing Snapshots and Schedules - BUI, CLI
- Renaming a Snapshot - BUI, CLI
- Editing a Snapshot Retention Policy - BUI, CLI

You can look at the contents of filesystem snapshots through the `.zfs/snapshot` filesystem directory. LUN snapshots cannot be accessed directly, though they can be used as a rollback target or as the source of a clone. To manage and access the `.zfs/snapshot` directory, use the following tasks:

- Making a Filesystem Snapshot Directory Visible - BUI, CLI
- Accessing a Hidden Filesystem Snapshot Directory (CLI)
- Accessing a Visible Filesystem Snapshot Directory (CLI)

You can use an existing snapshot to restore a filesystem or LUN to the exact state it was in when the snapshot was taken. To roll back a filesystem, LUN, or project to an existing snapshot, use the following task: Rolling Back to a Snapshot - BUI, CLI.

To destroy snapshots, use the following task: Destroying a Snapshot - BUI, CLI.

Snapshot Space Management

Snapshots present an interesting dilemma for space management. They represent the set of physical blocks referenced by a share at a given point in time. Initially, this snapshot consumes no additional space. But as new data is overwritten in the new share, the blocks in the active share will only contain the new data, and older blocks will be "held" by the most recent (and possibly older) snapshots. Gradually, snapshots can consume additional space as the content diverges in the active share. If you take a snapshot of a filesystem of any given size, and re-write 100 percent of the data within the filesystem, you must maintain references to twice the data that was originally in the filesystem.

Each snapshot has two associated space statistics: unique space and referenced space. The amount of referenced space is the total space consumed by the filesystem at the time the snapshot was taken. It represents the theoretical maximum size of the snapshot should it remain the sole reference to all data blocks. The unique space indicates the amount of physical space referenced only by the current snapshot. When a snapshot is destroyed, the unique space is made available to the rest of the pool.

Note that the amount of space consumed by all snapshots is not equivalent to the sum of unique space across all snapshots. With a share and a single snapshot, all blocks must be referenced by one or both of the snapshot or the share. With multiple snapshots, however, it's possible for a block to be referenced by some subset of snapshots, and not any particular snapshot. For example, if a file is created, two snapshots X and Y are taken, the file is deleted, and another snapshot Z is taken, the blocks within the file are held by X and Y, but not by Z. In this case, destroying Z will not free up the space, but destroying both X and Y will. Because of this, destroying any snapshot can affect the unique space referenced by neighboring snapshots, though the total amount of space consumed by snapshots will always decrease.

The total size of a project or share always accounts for space consumed by all snapshots, though the usage breakdown is also available. Quotas and reservations can be set at the project level to enforce physical constraints across this total space. In addition, quotas and reservations can be set at the filesystem level, and these settings can apply to only referenced data or total data.

Whether or not quotas and reservations should be applied to referenced data or total physical data depends on the administrative environment. If users are not in control of their snapshots (that is, an automatic snapshot schedule is set for them), then quotas should typically not include snapshots in the calculation. Otherwise, the user may run out of space but be confused when files cannot be deleted. Without an understanding of snapshots or means to manage those snapshots, it is possible for such a situation to be unrecoverable without administrator intervention. In this scenario, the snapshots represent an overhead cost that is factored into operation of the system in order to provide backup capabilities. On the other hand, there are environments where users are billed according to their physical space requirements, and snapshots represent a choice by the user to provide some level of backup that meets their requirements given the churn rate of their dataset. In these environments, it makes more sense to enforce quotas based on total physical data, including snapshots. The users understand the cost of snapshots, and can be provided a means to actively manage them (as through dedicated roles on the appliance).

Related Topics

- Space Management for Shares
- Managing User-Generated Snapshots

Taking a Snapshot (BUI)

Use the following procedure to take a manual snapshot of a filesystem, LUN, or project, and to specify the retention policy setting. A manually set retention policy, which is especially beneficial for legal holds, has no specified duration, and the snapshot cannot be deleted until the `off` option is manually set. A snapshot with a set retention policy within a share also protects the share and project containing the snapshot from deletion. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. Snapshot retention holds are preserved when moving the snapshot to another system via remote replication, NDMP (zfs format), or cloud snapshot backup (zfs format). However, retention holds cannot be added to original remote replication snapshots nor original NDMP snapshots.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to take a manual snapshot and to set a retention policy:

- Take a manual snapshot: `takeSnap`
- Set a retention policy: `retainSnap`

To release a hold so a snapshot can be deleted, a user with the `releaseSnapRetention` authorization must set the retention policy to `off`. For information on editing authorizations for a role, see Editing Authorizations for a Role (BUI). For information on editing the retention policy, see Editing a Snapshot Retention Policy (BUI).

To schedule automatic snapshots at regular intervals and to specify the retention policy setting, see Scheduling Snapshots (BUI).

1. Go to the share or project you want to snapshot.
   1. To take a snapshot of a filesystem: From the **Shares** menu, select **Shares**.
   2. To take a snapshot of a LUN: From the **Shares** menu, select **Shares**, then **LUNs**.
   3. To take a snapshot of a project: From the **Shares** menu, select **Projects**.
2. Hover over the share or project, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. Click the add icon ⊕ next to **Snapshots**.
5. Type a name for the snapshot. Select a retention policy option:
   - **Off** - No retention policy applies to this snapshot. This is the default setting.
   - **Unlocked -** This option can be selected by a user with the `retainSnap` authorization. It sets a retention hold for this snapshot for an unspecified duration, and the snapshot cannot be deleted.

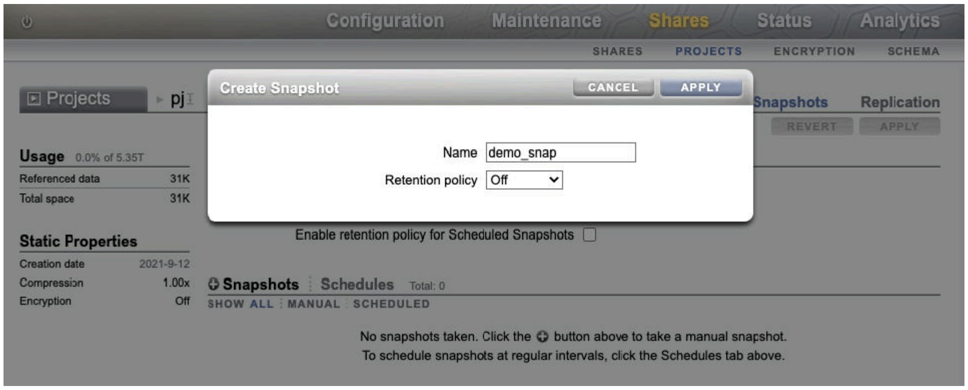

Name demo_snap
Retention policy Off

Usage 0.0% of 5.35T
Referenced data 31K
Total space 31K

Static Properties
Creation date 2021-9-12
Compression 1.00x
Encryption Off

Enable retention policy for Scheduled Snapshots □

No snapshots taken. Click the ⊕ button above to take a manual snapshot.
To schedule snapshots at regular intervals, click the Schedules tab above.

6. Click **APPLY**.

Taking a Snapshot (CLI)

Use the following procedure to take a manual snapshot of a filesystem, LUN, or project, and to specify the retention policy setting. A manually set retention policy, which is especially beneficial for legal holds, has no specified duration, and the snapshot cannot be deleted until the `off` option is manually set. A snapshot with a set retention policy within a share also protects the share and project containing the snapshot. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. Snapshot retention holds are preserved when moving the snapshot to another system via remote replication, NDMP (zfs format), or cloud snapshot backup (zfs format). However, retention holds cannot be added to original remote replication snapshots nor original NDMP snapshots.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to take a manual snapshot and to set a retention policy:

- Take a manual snapshot: `takeSnap`
- Set a retention policy: `retainSnap`

To release a hold so a snapshot can be deleted, a user with the releaseSnapRetention authorization must set the retention policy to off. For information on editing authorizations for a role, see Editing Authorizations for a Role (BUI). For information on editing the retention policy, see Editing a Snapshot Retention Policy (CLI).

To schedule automatic snapshots at regular intervals and to specify the retention policy setting, see Scheduling Snapshots (CLI).

1. Go to the share or project you want to snapshot.
   a. To take a snapshot of a project, go to shares and select the project.

      ```
      hostname:shares> select myproject
      hostname:shares myproject>
      ```

   b. To take a snapshot of a filesystem or LUN, go to shares and select the project containing the share, then select the share.

      ```
      hostname:shares> select myproject
      hostname:shares myproject> select demo_share
      hostname:shares myproject/demo_share>
      ```

2. Enter snapshots.

   ```
   hostname:shares myproject/demo_share> snapshots
   ```

3. Use the snapshot command followed by the name you want to give the new snapshot. To set a retention hold, use the r option after the snapshot command.

   The retention hold option can be set by a user with the retainSnap authorization. It sets a retention policy for this snapshot for an unspecified hold duration, the retentionpolicy property is set to unlocked, and the snapshot cannot be deleted.

   ```
   hostname:shares myproject/demo_share snapshots> snapshot -r demo_snap
   ```

Scheduling Snapshots (BUI)

Use the following procedure to configure automatic snapshots of a filesystem, LUN, or project and set a retention policy for those snapshots.

Automatic snapshots can be set on a project or a share, but not both. Otherwise, overlapping schedules and retention policies would make it impossible to guarantee both schedules. Removing an interval, or changing its retention policy, will immediately destroy any automatic snapshots not covered by the new schedule, unless the snapshots are under a retention hold. Automatic snapshots with clones are ignored.

Automatic snapshots can be taken half-hourly, hourly, daily, weekly, or monthly and are named .auto[-<snaplabel>]-<timestamp>. In the **Creation** column of the **Snapshots** list, times are displayed in the local (client browser) time zone. However, times are stored and executed in UTC format, without regard to such conventions as daylight saving time. For example, a snapshot scheduled for 10:00 a.m. PST (UTC-8) is stored and executed at 18:00 UTC, and this is the time that will appear as the timestamp in the snapshot name.

Older versions of the software allowed for automatic snapshots at the frequency of a minute. To help users avoid placing undue stress on the system, this feature was removed with the 2010.Q3 release. If the software is rolled back, existing minutes will be preserved. Previous instances will expire according to the existing schedule, but no new snapshots will be taken. An alert will be posted if a share or project with this frequency is found.

Automatic snapshots can be kept forever (except for half-hourly and hourly snapshots, which are capped at 48 and 24, respectively), or they can be limited to a certain number. When the number of snapshots exceeds the number for the **Keep at most** property, the oldest snapshots are deleted first. This is known as the "keep-at-most" scheme.

A retention hold can be placed on automatic snapshots to prevent deletion by locking them, which is especially beneficial for meeting internal retention and regulatory compliance needs. Because a retention policy is set at the parent level, the project and its shares are equally protected from deletion. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. Snapshot retention holds are preserved when moving the snapshot to another system via remote replication, NDMP (zfs format), or cloud snapshot backup (zfs format). However, retention holds cannot be added to original remote replication snapshots nor original NDMP snapshots.

When the threshold is met for the **Retention** property, the retention hold is set to `off` for excessive snapshots, and they are automatically deleted using the keep-at-most scheme. Also, a retention hold cannot be modified nor removed. Therefore, a snapshot schedule with a retention hold cannot be deleted until it is the oldest snapshot in the keep-at-most scheme. For a snapshot within a share, its share and project also cannot be deleted early.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to configure an automatic snapshot:

- Schedule an automatic snapshot: `scheduleSnap`
- Set a retention hold: `scheduleLockedSnap`

For information on editing authorizations for a role, see Editing Authorizations for a Role (BUI). For information on editing the retention hold value, see Editing a Snapshot Retention Policy (BUI).

1. Go to the project.
2. To set a retention policy for the project and its shares, select the check box for **Enable retention policy for Scheduled Snapshots**.
3. To schedule snapshots for a share, go to the appropriate share:
   a. For a filesystem: From the **Shares** menu, select **Shares**.
   b. For a LUN: From the **Shares** menu, select **Shares**, then **LUNs**.
4. To set a retention policy for the share, select the check box for **Enable retention policy for Scheduled Snapshots**.

   If you set the retention policy at the project level, you cannot set it at the share level because it is already set.
5. Hover over the share or project, and click the edit icon ✎ .
6. Click the **Snapshots** tab.
7. Click **Schedules**.
8. Click the add icon ⊕ next to **Schedules**.
9. Set each field appropriately.
   a. Set the frequency to **half-hourly**, **hourly**, **daily**, **weekly**, or **monthly** to indicate how often the snapshot is automatically taken.
   b. Set the precise time the snapshot is automatically taken.

      For half-hourly or hourly snapshots, you can choose how many minutes after the half-hour or hour the snapshot is taken. For daily snapshots, you can choose the hour and minute the snapshot is taken, and for weekly or monthly snapshots, you can specify the day, hour, and minute.

c. To not set a retention policy, clear the **Keep at most** check box. To set a retention policy, keep the check box selected, and set the **Keep at most** property to the number of snapshots to keep before older snapshots are automatically deleted.

   If multiple snapshot schedules end at the same time, only the snapshot with the strictest retention hold is generated.

d. To set a retention hold, set the **Retention** property to **Locked**, and specify the number of snapshots that should be locked and, thus, protected from deletion.

   The number for the **Retention** property must be the same or smaller than the number for the **Keep at most** property. The default for the **Retention** property is **Off**.

10. Click **APPLY**.

Scheduling Snapshots (CLI)

Use the following procedure to configure automatic snapshots of a share and set a retention policy for those snapshots.

Automatic snapshots can be set on a project or a share, but not both. Otherwise, overlapping schedules and retention policies would make it impossible to guarantee both schedules. Removing an interval, or changing its retention policy, will immediately destroy any automatic snapshots not covered by the new schedule, unless the snapshots are under a retention hold. Automatic snapshots with clones are ignored.

Automatic snapshots can be taken half-hourly, hourly, daily, weekly, or monthly and are named `.auto[-<snaplabel>]-<timestamp>`. Snapshot creation times are stored and executed in UTC format, without regard to such conventions as daylight saving time. For example, a snapshot scheduled for 10:00 a.m. PST (UTC-8) is stored and executed at 18:00 UTC, and this is the time that will appear as the timestamp in the snapshot name.

Older versions of the software allowed for automatic snapshots at the frequency of a minute. To help users avoid placing undue stress on the system, this feature was removed with the 2010.Q3 release. If the software is rolled back, existing minutes will be preserved. Previous instances will expire according to the existing schedule, but no new snapshots will be taken. An alert will be posted if a share or project with this frequency is found.

Automatic snapshots can be kept forever (except for half-hourly and hourly snapshots, which are capped at 48 and 24, respectively), or they can be limited to a certain number. When the number of snapshots exceeds the number for the `keep` property, the oldest snapshots are deleted first. This is known as the "keep-at-most" scheme.

A retention hold can be placed on automatic snapshots to prevent deletion by locking them, which is especially beneficial for meeting internal retention and regulatory compliance needs. Because a retention policy is set at the parent level, the project and its shares are equally protected from deletion. However, filesystems, LUNs, and other snapshots within the share can be modified or deleted. Snapshot retention holds are preserved when moving the snapshot to another system via remote replication, NDMP (zfs format), or cloud snapshot backup (zfs format). However, retention holds cannot be added to original remote replication snapshots nor original NDMP snapshots.

When the threshold is met for the `retentionhold` property, the retention hold is set to `off` for excessive snapshots, and they are automatically deleted using the keep-at-most scheme. Also, a retention hold cannot be modified nor removed. Therefore, a snapshot schedule with a retention hold cannot be deleted until it is the oldest snapshot in the keep-at-most scheme. For a snapshot within a share, its share and project also cannot be deleted early.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to configure an automatic snapshot:

- Schedule an automatic snapshot: `scheduleSnap`
- Set a retention hold: `scheduleLockedSnap`

For information on editing authorizations for a role, see Editing Authorizations for a Role (CLI). For information on editing the retention hold value, see Editing a Snapshot Retention Policy (CLI).

1. Go to shares and select the project or share you want to snapshot.

   ```
   hostname:> shares select myproject
   hostname:shares myproject> select demo_share
   hostname:shares myproject/demo_share>
   ```

2. To enable the retention policy for automatic snapshots within the project or share, set property `snapret_enabled` to `true`.

   ```
   hostname:shares myproject/demo_share> set snapret_enabled=true
   ```

3. Enter `snapshots automatic`.

   ```
   hostname:shares myproject/demo_share> snapshots automatic
   hostname:shares myproject/demo_share snapshots automatic>
   ```

4. Enter the `create` command to enter an uncommitted schedule context.

   ```
   hostname:shares myproject/demo_share snapshots automatic> create
   hostname:shares myproject/demo_share snapshots automatic (uncommitted)>
   ```

5. Use the `set` command to set each field appropriately.

   a. Set the frequency to `halfhour`, `hour`, `day`, `week`, or `month` to indicate how often the snapshot is automatically taken.

   b. Set the day, hour, and minute to specify the precise time the snapshot is automatically taken.

      For half-hourly or hourly snapshots, you can choose how many minutes after the half-hour or hour the snapshot is taken. For daily snapshots, you can choose the hour and minute the snapshot is taken, and for weekly or monthly snapshots, you can specify the day, hour, and minute.

   c. To set a retention policy, set the `retentionpolicy` property to `locked`.

      The default for the `retentionpolicy` property is `off`.

   d. To set the retention policy value, set the keep property to the number of snapshots to keep before older snapshots are automatically deleted.

      If multiple snapshot schedules end at the same time, only the snapshot with the strictest retention hold is generated.

   e. To set a retention hold, set the `retentionhold` property to the number of snapshots that should be locked and, thus, protected from deletion.

      The number for the `retentionhold` property must be the same or smaller than the number for the `keep` property.

   ```
   hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set
   frequency=day
                        frequency = day (uncommitted)
   ```

```
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set hour=14
                          hour = 14 (uncommitted)
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set
minute=30
                        minute = 30 (uncommitted)
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set keep=7
                          keep = 7 (uncommitted)
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set
retentionhold=3
                 retentionhold = 3 (uncommitted)
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> set
retentionpolicy=locked
                 retentionhold = 3 (uncommitted)
```

You can use the `get` command to view the current uncommitted settings.

```
hostname:shares myproject/demo_share snapshots automatic (uncommitted)> get
                     frequency = day (uncommitted)
                           day = (unset)
                          hour = 14 (uncommitted)
                        minute = 30 (uncommitted)
                          keep = 7 (uncommitted)
                 retentionhold = 3 (uncommitted)
               retentionpolicy = locked (uncommitted)
```

6. Enter `commit` to commit the changes, and create the automatic snapshot schedule.

   ```
   hostname:shares myproject/demo_share snapshots automatic (uncommitted)> commit
   ```

   You can use the `list` command to view the new schedule.

   ```
   hostname:shares myproject/demo_share snapshots automatic> list
   NAME                 FREQUENCY       DAY           HH:MM KEEP
   automatic-000        day             -             14:30    7
   ```

7. Enter `done` to finish.

   ```
   hostname:shares myproject/demo_share snapshots automatic> done
   hostname:shares myproject/demo_share snapshots>
   ```

Setting a Scheduled Snapshot Label (BUI)

Use the following procedure to set a label for scheduled snapshots of a filesystem, LUN, or project.

This optional property appends a user-defined label to each scheduled snapshot and is blank by default. The label can either be set for an individual share, or it can be set for a project and inherited by its shares, but not both.

Snapshot labels can help identify the project or share for which a snapshot was taken. For example, `project1:share1` could indicate a scheduled snapshot taken on `share1` within `project1`. Labels can be up to 35 alphanumeric characters and can include special characters `_-.:.`

1. Go to the share or project for which you want to set the scheduled snapshot label.
   a. To set a label for a filesystem: From the **Shares** menu, select **Shares**.
   b. To set a label for a LUN: From the **Shares** menu, select **Shares**, then **LUNs**.
   c. To set a label for a project: From the **Shares** menu, select **Projects**.

2. Hover over the appropriate share or project, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. Under **Properties**, type the label you want to set into the **Scheduled Snapshot Label** field.
5. Click **APPLY** to save the change.

   This label will be included in the name of each scheduled snapshot taken from now on. The label will appear before the timestamp, so that the snapshot name is `.auto-<snaplabel>-<timestamp>`.

Setting a Scheduled Snapshot Label (CLI)

Use the following procedure to set a label for scheduled snapshots of a filesystem, LUN, or project.

This optional property appends a user-defined label to each scheduled snapshot and is blank by default. The label can either be set for an individual share, or it can be set for a project and inherited by its shares, but not both.

Snapshot labels can help identify the project or share for which a snapshot was taken. For example, `project1:share1` could indicate a scheduled snapshot taken on `share1` within `project1`. Labels can be up to 35 alphanumeric characters and can include special characters **_-.:.**

1. Go to `shares` and select the filesystem, LUN, or project for which you want to set the label.

```
hostname:shares myproject> select demo_share
hostname:shares myproject/demo_share>
```

2. Use the `set snaplabel` command to create a scheduled snapshot label.

```
hostname:shares myproject/demo_share> set snaplabel=myproject:demo_share
```

Viewing Snapshots and Schedules (BUI)

Use the following procedure to view the snapshots and automatic snapshot schedules of a particular filesystem, LUN, or project.

1. Go to the share or project.
   a. To view snapshots and snapshot schedules of a filesystem: From the **Shares** menu, select **Shares**.
   b. To view snapshots and snapshot schedules of a LUN: From the **Shares** menu, select **Shares**, then **LUNs**.
   c. To view snapshots and snapshot schedules of a project: From the **Shares** menu, select **Projects**.
2. Hover over the share or project, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. View the snapshots or snapshot schedules.
   a. View the snapshots of that share under **Snapshots**, and optionally select **MANUAL** or **SCHEDULED** to view only manual or only scheduled snapshots.

Click the edit icon ✎ for a snapshot to view its details: name (**NAME**), date and time of creation (**CREATION**), amount of unique space used by the snapshot (**UNIQUE**), total amount of space referenced by the snapshot (**TOTAL**), and number of clones of the snapshot (**CLONES**).

b. Click **Schedules** to view the automatic snapshot schedules for that share.

   For each schedule, you can see the frequency that a snapshot is taken, the precise day and time at which it is taken, how many snapshots are kept before old ones are deleted, and how many snapshots have a retention hold.

Viewing Snapshots and Schedules (CLI)

Use the following procedure to view the snapshots and automatic snapshot schedules of a particular filesystem, LUN, or project.

1. Go to `shares` and select the project or share.

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
hostname:shares myproject/demo_share>
```

2. Enter `snapshots`.

```
hostname:shares myproject/demo_share> snapshots
hostname:shares myproject/demo_share snapshots>
```

3. View the snapshots or snapshot schedules using the appropriate commands.

   a. Enter `list` to view a list of the snapshots of this share or project.

```
hostname:shares myproject/demo_share snapshots> list
demo_snap1
demo_snap2
hostname:shares myproject/demo_share snapshots>
```

   You can select a snapshot, and use the `list` command to see the following properties:

   - `creation` - the date and time of creation of the snapshot in UTC format
   - `numclones` - the number of clones of the snapshot
   - `isauto` - whether the snapshot was created manually (`false`) or with an automatic snapshot schedule (`true`)
   - `retentionpolicy` - the retention policy setting of the snapshot
   - `pool` - what storage pool the snapshot is in
   - `canonical_name` - the location of the snapshot
   - `shadowsnap` - whether the snapshot was taken during shadow migration (`true`) or not (`false`)
   - `space_unique` - the amount of unique space the snapshot uses
   - `space_data` - the total amount of space the snapshot references

```
hostname:shares myproject/demo_share snapshots> select demo_snap1
hostname:shares myproject/demo_share snapshots demo_snap1> list
Properties:
           creation = Thu Nov 12 2021 20:19:49 GMT+0000(UTC)
          numclones = 1
```

```
             isauto = false
    retentionpolicy = off
               pool = pool1
     canonical_name = pool1/local/myproject/demo_share@demo_snap1
         shadowsnap = false
       space_unique = 0
         space_data = 31K
```

b. Enter `automatic`, and use the `list` command to view a list of the automatic snapshot schedules of this share or project.

```
hostname:shares myproject/demo_share snapshots> automatic
hostname:shares myproject/demo_share snapshots automatic> list
Properties:
          convert = false

Automatics:

NAME           FREQUENCY   DAY   HH:MM KEEP
automatic-000  day         -     00:00    4
automatic-001  month       01    00:00   12
```

Editing a Snapshot Retention Policy (BUI)

Use the following procedure to edit a snapshot retention policy for a filesystem, LUN, or project. Snapshot retention policies are included in both manual snapshots and in automatic snapshot schedules.

After an automatic snapshot with a retention hold has been generated, it is not possible to delete or modify its schedule until all snapshots with a retention hold have expired.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to make changes to a retention policy:

- Modify a manual snapshot to a stricter hold type (`off` to `unlocked`): `retainSnap`
- Modify a manual snapshot to a less strict hold type (`unlocked` to `off`): `releaseSnapRetention`
- Modify an automatic snapshot schedule: `scheduleSnap`
- Modify the retention hold value for an automatic snapshot: `releaseSnapRetention`

For information on editing authorizations for a role, see Editing Authorizations for a Role (BUI). To delete an automatic snapshot schedule, see Removing a Snapshot Schedule (BUI).

1. Go to the appropriate project or share.
   a. If the schedule applies to a project: From the **Shares** menu, select **Projects**.
   b. If the schedule applies to a filesystem: From the **Shares** menu, select **Shares**.
   c. If the schedule applies to a LUN: From the **Shares** menu, select **Shares**, then **LUNs**.
2. Hover over the appropriate project or share, and click the edit icon ✎ .
3. Select the check box for **Enable retention policy for Scheduled Snapshots** to set a retention policy. To not set a retention policy, clear the check box.
4. Click the **Snapshots** tab.

5. For a manual snapshot and to change the retention policy, change the retention hold type **off** to **unlocked**, or **unlocked** to **off**, as appropriate.

   If the snapshot is a project snapshot, this setting also applies to all of its shares.

6. For an automatic snapshot, click **Schedules**.

7. Clear the **Keep at most** check box to not set a retention policy. To change the retention policy value, set the **Keep at most** property to the number of snapshots to keep before older snapshots are automatically deleted.

   If the schedule includes snapshots with a retention hold, the **Keep at most** property can only be changed to a higher value. If no automatic snapshots have been generated with this schedule, or none of the snapshots have a retention hold, the **Keep at most** property can be set to a lower value. If the snapshot is a project snapshot, this setting also applies to all of its shares.

8. To change the retention hold if the retention hold type is **off**, set the **Retention** property to **locked**, and edit the retention hold value for the schedule.

   The following guidelines apply to automatic snapshots:

   - If the retention hold type is **locked**, you cannot edit the **Retention** property. When all locked snapshots for this schedule have exceeded the **Keep at most** property value, the retention hold value changes from **locked** to **off**.
   - The number for the **Retention** property must be the same or smaller than the number for the **Keep at most** property.
   - If automatic snapshots containing a retention hold have been generated with this schedule, the retention hold must be set to a higher value to prevent early lock removal, but not higher than the **Keep at most** property.
   - If no automatic snapshots have been generated with this schedule, the retention hold can be set to a lower value.
   - If the snapshot is a project snapshot, this setting also applies to all of its shares.

Editing a Snapshot Retention Policy (CLI)

Use the following procedure to edit a snapshot retention policy for a filesystem, LUN, or project. Snapshot retention policies are included in both manual snapshots and in automatic snapshot schedules.

After an automatic snapshot with a retention hold has been generated, it is not possible to delete or modify its schedule until all snapshots with a retention hold have expired.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to make changes to a retention policy:

- Modify a manual snapshot to a stricter hold type (`off` to `unlocked`): `retainSnap`
- Modify a manual snapshot to a less strict hold type (`unlocked` to `off`): `releaseSnapRetention`
- Modify an automatic snapshot schedule: `scheduleSnap`
- Modify the retention hold value for an automatic snapshot: `releaseSnapRetention`

For information on editing authorizations for a role, see Editing Authorizations for a Role (CLI). To delete an automatic snapshot schedule, see Removing a Snapshot Schedule (CLI).

1. Go to shares and select the project or share.

   ```
   hostname:> shares select myproject
   hostname:shares myproject> select demo_share
   hostname:shares myproject/demo_share>
   ```

2. Enter the appropriate node:

   - Manual snapshot: Enter `snapshots`
   - Automatic snapshot: Enter `snapshots automatic`

3. Select the snapshot to edit.

4. Set property `snapret_enabled` to `true` to set a retention policy. To not set a retention policy, set `snapret_enabled` to `false`.

   ```
   hostname:shares myproject/demo_share> set snapret_enabled=true
   ```

5. For a manual snapshot and to change the retention policy, enter the `set retentionpolicy` command followed by the appropriate retention policy option: either `off` to `unlocked`, or `unlocked` to `off`.

   If the snapshot is a project snapshot, this setting also applies to all of its shares.

6. For an automatic snapshot and to change the number of snapshots retained before older snapshots are automatically deleted, set the keep property to a new value.

   If the schedule includes snapshots with a retention hold, the `keep` property can only be changed to a higher value. If no automatic snapshots have been generated with this schedule, or none of the snapshots have a retention hold, the `keep` property can be set to a lower value. If the snapshot is a project snapshot, this setting also applies to all of its shares.

7. For an automatic snapshot and to change the retention hold if the retention hold type is `off`, set the `retentionpolicy` property to `locked`, and edit the `retentionhold` property for the schedule.

   The following guidelines apply to automatic snapshots:

   - If the retention hold type is `locked`, you cannot edit the `retentionpolicy` property. When all locked snapshots for this schedule have exceeded the `keep` property value, the retention hold value changes from `locked` to `off`.
   - The number for the `retentionpolicy` property must be the same or smaller than the number for the `keep` property.
   - If automatic snapshots containing a retention hold have been generated with this schedule, the `retentionhold` property must be set to a higher value to prevent early lock removal, but not higher than the `keep` property.
   - If no automatic snapshots have been generated with this schedule, the retention hold can be set to a lower value.
   - If the snapshot is a project snapshot, this setting also applies to all of its shares.

8. Enter `commit` to save the changes.

Removing a Snapshot Schedule (BUI)

Use the following procedure to delete an automatic snapshot schedule for a filesystem, LUN, or project.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to remove a snapshot schedule:

- Delete an automatic snapshot schedule without a retention hold: `scheduleSnap`
- Delete an automatic snapshot schedule with a retention hold: `releaseSnapRetention`

For information on editing authorizations for a role, see Editing Authorizations for a Role (BUI).

1. Go to the appropriate project or share.
    a. If the schedule applies to a project: From the **Shares** menu, select **Projects**.
    b. If the schedule applies to a filesystem: From the **Shares** menu, select **Shares**.
    c. If the schedule applies to a LUN: From the **Shares** menu, select **Shares**, then **LUNs**.
2. Hover over the appropriate project or share, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. Click **Schedules**.
5. Hover over the schedule you want to remove, and click the remove icon 🗑 .

   An error message might warn that existing automatic snapshots could be destroyed. If the snapshot or its children are actively changing, an error message indicates that the snapshot schedule cannot be removed. Also, if the schedule contains locked automatic snapshots, the schedule cannot be removed until the retention holds expire. If the automatic snapshot schedule has a retention hold but no snapshots have been generated, the schedule can be removed. If the snapshot is a project snapshot, the schedule will also be removed from its shares.
6. If you want to keep existing automatic snapshots, click **CONVERT** to convert them to manual snapshots. Otherwise, click **DISCARD** to destroy them. If the snapshots have a retention hold, **CONVERT** does not change the retention hold.

Removing a Snapshot Schedule (CLI)

Use the following procedure to delete an automatic snapshot schedule for a filesystem, LUN, or project.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to remove a snapshot schedule:

- Delete an automatic snapshot schedule without a retention hold: `scheduleSnap`
- Delete an automatic snapshot schedule with a retention hold: `releaseSnapRetention`

For information on editing authorizations for a role, see Editing Authorizations for a Role (CLI).

1. Go to and select the project or share.

   ```
   hostname:> shares select myproject
   hostname:shares myproject> select demo_share
   ```

2. Enter `snapshots automatic`.

```
hostname:shares myproject/demo_share> snapshots automatic
hostname:shares myproject/demo_share snapshots automatic>
```

3. If you want to keep existing automatic snapshots, set the `convert` property to `true` to convert them to manual snapshots.

   If you keep this property as `convert = false`, automatic snapshots taken with this schedule will be discarded when you destroy the schedule. If the snapshots have a retention hold, setting the `convert` property to `true` does not change the retention hold.

```
hostname:shares myproject/demo_share snapshots automatic> set convert=true
                              convert = true
hostname:shares myproject/demo_share snapshots automatic> commit
```

4. Use the `destroy` command followed by the name of the automatic snapshot schedule you want to destroy.

   An error message might warn that existing automatic snapshots could be destroyed. If the snapshot or its children are actively changing, an error message indicates that the snapshot schedule cannot be removed. Also, if the schedule contains locked automatic snapshots, the schedule cannot be removed until the retention holds expire. If the automatic snapshot schedule has a retention hold but no snapshots have been generated, the schedule can be removed. If the snapshot is a project snapshot, the schedule will also be removed from its shares.

```
hostname:shares myproject/demo_share snapshots automatic> destroy automatic-000
This will destroy "automatic-000". Are you sure? (Y/N)
```

5. Type `Y` to confirm.

```
This will destroy "automatic-000". Are you sure? (Y/N) Y
```

Working With the File System Control Directory (`.zfs`)

A file system stores snapshots and clones in its `.zfs` control directory. Note that the control directory is not visible in directory listings by default. To make the control directory visible, see Making the File System Control Directory (`.zfs`) Visible (BUI) and Making the File System Control Directory (`.zfs`) Visible (CLI). Note that having the control directory hidden prevents backup software from inadvertently backing up snapshots in addition to new data.

The file system control directory contains the following `snapshot` and `clone` directories:

- `.zfs/snapshot`: Contains a list of all snapshots on the file system. You can access these read-only snapshots in the same way that you access other file-system data.
- `.zfs/clone`: Contains a list of all clones of the file system that have their `mountpoint` property set to the `clonedir` value. You can access the clones mounted here in the same way that you can access other file-system data. The clones are shared with the file system of which they are a clone.

  By default, a file system or clone is mounted at `/export/`*name*, where *name* is the name of the file system, snapshot, or clone.

To prepare a file system share to use the control directory, see Mounting a Clone File System in the Control Directory (BUI) and Mounting a Clone File System in the Control Directory (CLI).

To access the hidden and visible file system snapshot directories, see Accessing the Hidden File System Snapshot Directory (CLI) and Accessing the Visible File System Snapshot Directory (CLI).

Making the File System Control Directory (.zfs) Visible (BUI)

> **Note:**
>
> Setting the .zfs directory to **Visible** might cause backup software to back up snapshots and clones in addition to live data.

1. From the **Shares** menu, select **Shares**.
2. Hover over the filesystem, and click the edit icon ✎.
3. Click the **Snapshots** tab.
4. Uncheck the **Inherit from project** box next to **Properties**, or click the lock icon 🔒 next to **.zfs visibility**.
5. Select **Visible** from the drop-down menu next to **.zfs visibility**.
6. Click **APPLY** to save the changes.

   To make the directory hidden again, return to this page, and select **Hidden** from the drop-down menu, then click **APPLY**.

Making the File System Control Directory (.zfs) Visible (CLI)

> **Note:**
>
> Setting the .zfs directory to visible may cause backup software to back up snapshots and clones in addition to live data.

1. Go to and select the file system share.

   ```
   hostname:> shares select myproject
   hostname:shares myproject> select demo_share
   hostname:shares myproject/demo_share>
   ```

2. Use the set snapdir command to set the file system snapshot directory to visible.

   ```
   hostname:shares myproject/demo_share> set snapdir=visible
           snapdir=visible(uncommitted)
   ```

3. Enter commit to save the change.

   ```
   hostname:shares myproject/demo_share> commit
   ```

4. To make the directory hidden again, return to this context and use the set snapdir command to set the directory to hidden, then enter commit to save the change.

   ```
   hostname:shares myproject/demo_share> set snapdir=hidden
           snapdir=hidden(uncommitted)
   hostname:shares myproject/demo_share> commit
   ```

Related Topics

- Mounting a Clone File System in the Control Directory (BUI)
- Mounting a Clone File System in the Control Directory (CLI)

- Making the File System Control Directory (.zfs) Visible (BUI)
- Accessing the Visible File System Snapshot Directory (CLI)
- Accessing the Hidden File System Snapshot Directory (CLI)

Mounting a Clone File System in the Control Directory (BUI)

This task enables you to mount a clone file system in the control directory. When you perform this task, the specified clone file system is unshared and unmounted from the current mount point. Then, the clone file system is mounted in the source file system's control directory and is shared only with its source file system.

1. From the **Shares** menu, select **Shares**.
2. Hover over the file system snapshot and click the edit icon ✎.
3. Click the lock icon 🔒 next to **Mountpoint**.
4. Click the **use 'clonedir'** box.

   When you click this box, the clone appears in the `.zfs/clone` directory. Note that this action does not affect the visibility of the control directory.
5. Click **APPLY** to save the changes.

You can stop using the control directory for the file system share. Uncheck the **use 'clonedir'** box, specify the mount point to use in the **mountpoint** field, such as `/export/demo_share`, and click **APPLY** to save the changes.

Mounting a Clone File System in the Control Directory (CLI)

This task enables you to mount a clone file system in the control directory. When you perform this task, the specified clone file system is unshared and unmounted from the current mount point. Then, the clone file system is mounted in the source file system's control directory and is shared only with its source file system.

1. To configure a control directory for the file system share, set the `mountpoint` property value to `clonedir`.

   ```
   hostname:shares myproject/demo_share> set mountpoint=clonedir
       mountpoint=clonedir(uncommitted)
   ```

   When you run this command, the clone appears in the `.zfs/clone` directory. Note that this action does not affect the visibility of the control directory.
2. Commit the change.

   ```
   hostname:shares myproject/demo_share> commit
   ```

To stop using the control directory for the file system share, set the `mountpoint` property value to the mount point, such as `/export/demo_share`.

```
hostname:shares myproject/demo_share> set mountpoint=/export/demo_share
    mountpoint=/export/demo_share(uncommitted)
hostname:shares myproject/demo_share> commit
```

Accessing the Hidden File System Snapshot Directory (CLI)

Use the following procedure to access file system snapshots over data protocols at `.zfs/snapshot` in the root of your file system.

1. In a terminal window, go to the directory where you mounted the share.
2. View contents of the .zfs/snapshot directory.

   You can list snapshots of this file system and can view the contents of each snapshot.

Example 6-1 Accessing .zfs/snapshot

The following example shows that two file-system snapshots called demo_snap1 and demo_snap2 are mounted at /mnt/demo. The ls -l /mnt/demo/.zfs/snapshot/demo_snap1 command shows that the demo_snap1 snapshot contains three files called file1, file2 and file3.

```
$ ls -1 /mnt/demo
/mnt/demo
$ ls -1 /mnt/demo/.zfs/snapshot
demo_snap1
demo_snap2
$ ls -1 /mnt/demo/.zfs/snapshot/demo_snap1
file1
file2
file3
```

Related Topics

- Mounting a Clone File System in the Control Directory (BUI)
- Mounting a Clone File System in the Control Directory (CLI)
- Making the File System Control Directory (.zfs) Visible (BUI)
- Making the File System Control Directory (.zfs) Visible (CLI)
- Accessing the Visible File System Snapshot Directory (CLI)

Accessing the Visible File System Snapshot Directory (CLI)

Use the following procedure to access file system snapshots in the .zfs/snapshot directory after making it visible.

Before You Begin

Set the .zfs directory to visible as described in Making a Filesystem Snapshot Directory Visible (CLI).

1. In a terminal window, go to the directory where you mounted the share.
2. View contents of the .zfs/snapshot directory.

   You can list snapshots of this file system and can view the contents of each snapshot.

Renaming a Snapshot (BUI)

Use the following procedure to rename an existing manual snapshot. An automatic snapshot that has .auto, .rr, or .ndmp in its name cannot be renamed.

If a share snapshot that is part of a larger project snapshot is renamed, it will no longer be considered part of the same snapshot, and if any snapshot is renamed to have the same name as a snapshot in the parent project, it will be treated as part of the project snapshot.

Before You Begin

- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the projects and shares scope:
  - `renameSnap` - Allows renaming snapshots.
  - `rename` - Allows renaming projects and shares, including snapshot names.
- To add authorizations to a role, see Editing Authorizations for a Role (BUI).

1. Go to the share or project that contains the snapshot you want to rename.
   - To rename a filesystem snapshot: From the **Shares** menu, select **Shares**.
   - To rename a LUN snapshot: From the **Shares** menu, select **Shares**, then **LUNs**.
   - To rename a project snapshot: From the **Shares** menu, select **Projects**.
2. Hover over the share or project that contains the snapshot you want to rename, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. Under **Snapshots**, click the name of the snapshot you want to rename.

   The snapshot name changes to a text input box.
5. Type the new name for the snapshot.

   A name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters _ **- . :**
6. Press **Enter** to commit the change.

Related Topics

- Understanding Users and Roles
- User Authorizations

Renaming a Snapshot (CLI)

Use the following procedure to rename an existing manual snapshot. Automatic snapshots cannot be renamed.

If a share snapshot that is part of a larger project snapshot is renamed, it will no longer be considered part of the same snapshot, and if any snapshot is renamed to have the same name as a snapshot in the parent project, it will be treated as part of the project snapshot.

Before You Begin

- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the projects and shares scope:
  - `renameSnap` - Allows renaming snapshots.
  - `rename` - Allows renaming projects and shares, including snapshot names.
- To add authorizations to a role, see Editing Authorizations for a Role (CLI).

1. Go to shares and select the project, or select the project and then a share.

   ```
   hostname:> shares select myproject
   hostname:shares myproject> select demo_share
   ```
2. Enter `snapshots`.

   ```
   hostname:shares myproject/demo_share> snapshots
   ```

3. Enter `list` to view the list of snapshots for the project or share.

```
hostname:shares myproject/demo_share snapshots> list
demo_snap1
demo_snap2
```

4. To rename the snapshot, enter `rename` followed by the current snapshot name, a space, and then the new snapshot name.

   A name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters _ **- . :**

```
hostname:shares myproject/demo_share snapshots> rename demo_snap1 new_name
```

Related Topics

- Understanding Users and Roles
- User Authorizations

Rolling Back to a Snapshot (BUI)

Use the following procedure to roll back, or restore, a filesystem or LUN to an existing snapshot.

When a rollback occurs, any newer snapshots (and clones of newer snapshots) are destroyed, and the active data are reverted to the state when the snapshot was taken. Snapshots only include data, not properties, so any property settings changed since the snapshot was taken will remain. Changes to filesystem root directory access are lost during rollback.

The rollback is not allowed if the rollback would remove recent snapshots with a snapshot retention hold.

Note:

When the file retention feature, which is different from the snapshot retention feature, is used with snapshot rollback, certain restrictions can apply. Rollback can be performed on a filesystem with the privileged file retention policy, even when unexpired retained files exist. Filesystems with the mandatory file retention policy can never be rolled back, even when all retained files have expired. For more information, see File Retention Management.

Caution:

This procedure cannot be undone.

1. Go to the share or project that contains the snapshot you want to restore.
   - To restore a filesystem snapshot: From the **Shares** menu, select **Shares**.
   - To restore a LUN snapshot: From the **Shares** menu, select **Shares**, then **LUNs**.
2. Hover over the share that contains the snapshot you want to restore, and click the edit icon ✎ .
3. Click the **Snapshots** tab.

4. Hover over the snapshot you want to restore, click its rollback icon ⟲ and confirm your action.

Rolling Back to a Snapshot (CLI)

Use the following procedure to roll back, or restore, a filesystem or LUN to an existing snapshot.

Restoring a snapshot requires destroying any newer snapshots and their clones, and it reverts the share contents to what they were at the time the snapshot was taken. Property settings on the share are not affected, but changes to filesystem root directory access are lost during rollback.

The rollback is not allowed if the rollback would remove recent snapshots with a snapshot retention hold.

Note:

When the file retention feature, which is different from the snapshot retention feature, is used with snapshot rollback, certain restrictions can apply. Rollback can be performed on a filesystem with the privileged file retention policy, even when unexpired retained files exist. Filesystems with the mandatory file retention policy can never be rolled back, even when all retained files have expired. For more information, see File Retention Management.

Caution:

This procedure cannot be undone.

1. Go to and select the share that contains the snapshot you want to restore.

```
hostname:> shares select myproject
hostname:shares myproject> select demo_share
```

2. Enter `snapshots`.

```
hostname:shares myproject/demo_share> snapshots
```

3. Enter `list` to view the list of snapshots for the project or share.

```
hostname:shares myproject/demo_share snapshots> list
demo_snap1
demo_snap2
```

4. Select the snapshot you want to restore, then enter the `rollback` command.

```
hostname:shares myproject/demo_share snapshots> select demo_snap1
hostname:shares myproject/demo_share@demo_snap1> rollback
```

5. Type `Y` to confirm.

```
hostname:shares myproject/demo_share@demo_snap1> rollback
Rolling back will revert data to snapshot, destroying newer data. Active initiators
will be disconnected.

Continue? (Y/N)
hostname: shares myproject/demo_share@demo_snap1> Y
```

Destroying a Snapshot (BUI)

Use the following procedure to destroy a snapshot.

Before You Begin

- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the projects and shares scope:
  - `destroySnap` - Allows destroying snapshots.
  - `destroy` - Allows destroying projects and shares, including snapshot names.
- To add authorizations to a role, see Editing Authorizations for a Role (BUI).
- Before a manual snapshot with a retention hold can be destroyed, the hold type must be `off`. To modify a manual snapshot from `unlocked` to `off`, see Editing a Snapshot Retention Policy (BUI).

  To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

1. Go to the snapshot.
   - For a snapshot of a filesystem or LUN: From the **Shares** menu, select **Shares**, then **Filesystems** or **LUNs**, depending on whether the snapshot you want to destroy is of a filesystem or a LUN.
   - For a snapshot of a project: From the **Shares** menu, select **Projects**. Destroying a project snapshot also destroys its shares.
2. Hover over the appropriate share, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. Hover over the snapshot you want to destroy, and click the destroy icon 🗑 .

   A confirmation dialog box appears.

   If clones have been made of this snapshot, you are prompted with a list of the clones that will be affected. Destroying a snapshot also destroys any clones of that snapshot and descendants of those clones.

   If the snapshot or its children are actively changing, an error message indicates that the snapshot cannot be destroyed. Also, if a schedule contains locked automatic snapshots, the snapshot cannot be destroyed until the retention holds expire. If an automatic snapshot schedule has a retention hold but no snapshots have been generated, the snapshot can be destroyed.
5. Click **OK** to confirm.

Related Topics

- Understanding Users and Roles
- User Authorizations

Destroying a Snapshot (CLI)

Use the following procedure to destroy a snapshot.

Before You Begin

- To complete this procedure, you must have Super-User privileges or one of the following role authorizations within the projects and shares scope:
  - `destroySnap` - Allows users to only destroy snapshots.
  - `destroy` - Grants privileges to remove projects and shares, including snapshots.
- To add authorizations to a role, see Editing Authorizations for a Role (CLI).
- Before a manual snapshot with a retention hold can be destroyed, the hold type must be `off`. To modify a manual snapshot from `unlocked` to `off`, see Editing a Snapshot Retention Policy (CLI).

  To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

1. Go to shares and select the project, or select the project and then a share.

   Destroying a project snapshot also destroys its shares.

   ```
   hostname:> shares select myproject
   hostname:shares myproject> select demo_share
   ```

2. Enter `snapshots`.

   ```
   hostname:shares myproject/demo_share> snapshots
   ```

3. Enter `list` to view the list of snapshots for the project or share.

   ```
   hostname:shares myproject/demo_share snapshots> list
   demo_snap1
   demo_snap2
   ```

4. Use the `destroy` command to delete an individual snapshot using one of two methods:

   - Select the snapshot you want to delete, then enter `destroy`.

     ```
     hostname:shares myproject/demo_share snapshots> select demo_snap1
     hostname:shares myproject/demo_share@demo_snap1> destroy
     ```

   - Type `destroy` followed by the snapshot name.

     ```
     hostname:shares myproject/demo_share snapshots> destroy demo_snap1
     ```

   If clones have been made of this snapshot, you are prompted with a list of the clones that will be affected. Destroying a snapshot also destroys any clones of that snapshot and descendants of those clones.

   If the snapshot or its children are actively changing, an error message indicates that the snapshot cannot be destroyed. Also, if a schedule contains locked automatic snapshots, the snapshot cannot be destroyed until the retention holds expire. If an automatic snapshot schedule has a retention hold but no snapshots have been generated, the snapshot can be destroyed.

5. Type `Y` to confirm your action.

   ```
   This will destroy all data in "demo_snap1"! Are you sure? (Y/N) Y
   ```

Related Topics

- Understanding Users and Roles
- User Authorizations

Cloning a Snapshot (BUI)

> **Note:**
>
> Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Use the following procedure to make a clone of an existing snapshot of a filesystem or LUN.

> **Note:**
>
> Clones of projects are not supported.

1. Go to the share you want to clone.
   a. To clone a filesystem: From the **Shares** menu, select **Shares**.
   b. To clone a LUN: From the **Shares** menu, select **Shares**, then **LUNs**.
2. Hover over the share, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. Hover over the snapshot you want to clone, and click the clone icon ⊞ .

   A dialog box appears with settings and options for the new clone.
5. Set each field appropriately.
   a. From the **Project** drop-down menu, select the destination project.

      By default, the clone is created within the current project, but it can be created in a different project.
   b. Type a name for the clone.
   c. Optional: Click the lock icon 🔒 next to **Mountpoint**, and set a mountpoint for the clone.

      If you leave this field locked, the mountpoint for the clone will remain as the default `/export/<sharename>`.
   d. Optional: Click the lock icon 🔒 next to **Resource name**, and enter one of the following values:
      - **off** - SMB is disabled.
      - **on** - SMB is enabled, so you can share the clone over SMB. The name of the clone in SMB matches the name of the clone in Oracle ZFS Storage Appliance.
      - **<pick_a_name>** - SMB is enabled, so you can share the clone over SMB. The name of the clone in SMB is the name you specify here instead of the name of the clone in the appliance.

      If you leave this field locked, the **Resource name** property will inherit from the snapshot you are cloning.

e. Optional: Select the **Inherit key** check box or uncheck the check box, and select the keystore and the name of the encryption key that you want the clone to inherit.

   If the box is checked, the keystore and keyname of the clone will be that of the destination project.

   If the box is unchecked, the keystore and keyname of the clone will be that of the parent share. Alternatively, select a different keystore and keyname from the drop-down menu.

f. Optional: Check the **Retain Other Local Settings** check box to cause any inherited properties to be preserved as local settings in the new clone.

   This field determines whether inherited properties will come from the parent dataset or the destination project. By default, the box is unchecked, meaning that all inherited properties will come from the destination project for the new clone. If you check the box, all currently inherited properties will be preserved as local settings in the new clone.

   The clone will have the same retention hold setting as the original snapshot. To apply or remove a retention hold for the clone, make a snapshot of the clone and specify a new retention hold setting.

   To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

6. Click **APPLY** to confirm the settings, and to create the clone.

   The clone appears in the list of shares for the destination project you set. You can work with a clone just like any other share.

Related Topics

- To perform share operations on a clone, see Shares and Projects.
- To make a clone of a clone, see Cloning a Clone.
- To view all the clones of a particular snapshot, see Viewing Clones of a Snapshot (BUI).
- To determine the snapshot from which a clone was made, see Viewing a Clone Origin (BUI).
- To make a clone of a snapshot in a replication package, see Cloning a Snapshot in a Replication Package (BUI).

Cloning a Snapshot (CLI)

> **Note:**
>
> Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Use the following procedure to make a clone of an existing snapshot of a filesystem or LUN.

Note:

Clones of projects are not supported.

1. Go to the appropriate filesystem, LUN, or project and enter `snapshots`.

   ```
   hostname:shares myproject/demo_share> snapshots
   hostname:shares myproject/demo_share snapshots>
   ```

2. Select the snapshot you want to clone.

   ```
   hostname:shares myproject/demo_share snapshots> select snap1
   ```

3. Use the `clone` command, followed by the name of the local project in which you want to create the clone if different from the original project, and then the clone name.

   By default, the clone is created in the same project as the snapshot being cloned.

   ```
   hostname:shares myproject/demo_share@snap1> clone demo_clone
   ```

   You are placed into an uncommitted share context. From here, you can adjust properties as needed before committing the changes to create the clone.

   The clone will have the same retention hold setting as the original snapshot. To apply or remove a retention hold for the clone, make a snapshot of the clone and specify a new retention hold setting.

   To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

4. Use the `get` command to view properties.

   ```
   hostname:shares myproject/demo_clone (uncommitted clone)> get
                       aclinherit = restricted (inherited)
                          aclmode = discard (inherited)
                            atime = true (inherited)
                  casesensitivity = mixed
                         checksum = fletcher4 (inherited)
                      compression = off (inherited)
                            dedup = false (inherited)
                    compressratio = 100
                           copies = 1 (inherited)
                         creation = Mon Nov 15 2021 17:27:29 GMT+0000 (UTC)
                          logbias = latency (inherited)
                       mountpoint = /export/clone0 (inherited)
                    normalization = none
                            quota = 0
                       quota_snap = true
                         readonly = false (inherited)
                       recordsize = 128K (inherited)
                      reservation = 0
                 reservation_snap = true
                         rstchown = true (inherited)
                   secondarycache = all (inherited)
                           shadow = none
                           nbmand = false (inherited)
                         sharesmb = off (inherited)
                         sharenfs = off (inherited)
                          snapdir = hidden (inherited)
                         utf8only = true
                            vscan = false (inherited)
   ```

```
                    encryption = off
                    writelimit = unlimited (default)
                     readlimit = unlimited (default)
                     snaplabel =
                      sharedav = off (inherited)
                      shareftp = off (inherited)
                     sharesftp = off (inherited)
                     sharetftp = off (inherited)
              shareobjectstore = off (inherited)
                       shares3 = off (inherited)
                      shareoci = off (inherited)
                          pool = p0
                        origin = default/f0@snap0
                canonical_name = p0/local/default/clone0
            effectivereadlimit = unlimited
           effectivewritelimit = unlimited
                      exported = true (inherited)
                     nodestroy = false
                  maxblocksize = 1M (inherited)
                  lz4supported = true (inherited)
   acl_passthru_mode_supported = true (inherited)
                    space_data = 1K
              space_unused_res = 0
               space_snapshots = 0
               space_available = 5.35T
                   space_total = 1K
                      root_acl =
              root_permissions = 700
                     root_user = nobody
                    root_group = other
                   smbshareacl =
               snapret_enabled = false
```

5. Use the `set` command to adjust properties.

```
hostname:shares myproject/demo_clone (uncommitted clone)> set quota=10G
                         quota = 10G (uncommitted)
```

6. Use the `commit` command to commit the changes, and to create the clone.

```
hostname:shares myproject/demo_clone (uncommitted clone)> commit
hostname:shares myproject/demo_share@demo_clone>
```

Related Topics

- To perform share operations on a clone, see Shares and Projects.
- To make a clone of a clone, see Cloning a Clone.
- To view all the clones of a particular snapshot, see Viewing Clones of a Snapshot (CLI).
- To determine the snapshot from which a clone was made, see Viewing a Clone Origin (CLI).
- To make a clone of a snapshot in a replication package, see Cloning a Snapshot in a Replication Package (CLI).

Cloning a Clone

> **Note:**
>
> Cloning is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Use the following procedure to make a clone of an existing clone.

Once you have created a clone from a snapshot of a filesystem or LUN, you can work with that clone as with any other share. You can take a snapshot of the clone, and you can make a clone of that snapshot. You can continue this process to make clones of clones indefinitely.

1. Take a snapshot of the clone using one of these procedures:
   - Taking a Snapshot (BUI)
   - Taking a Snapshot (CLI)
2. Clone the snapshot using one of these procedures:
   - Cloning a Snapshot (BUI)
   - Cloning a Snapshot (CLI)

Viewing Clones of a Snapshot (BUI)

Use the following procedure to view a list of all clones created from a particular snapshot. These are also called the "dependent clones" of the snapshot.

1. From the **Shares** menu, select **Shares**, then either **Filesystems** or **LUNs**, depending on whether you want to view clones of a filesystem or a LUN.
2. Hover over the appropriate share, and click the edit icon ✎ .
3. Click the **Snapshots** tab.
4. Hover over the appropriate snapshot, and click **Show** under **Clones**.

   A window appears with a list of the snapshot's dependent clones and the projects in which they are located.

   If the **Show** link does not appear, the snapshot has no clones.
5. Click **OK** to close the window.

Viewing Clones of a Snapshot (CLI)

Use the following procedure to view a list of all clones created from a particular snapshot.

1. Go to and select the snapshot.

```
hostname:shares myproject/demo_share> snapshots
hostname:shares myproject/demo_share snapshots> select snap1
hostname:shares myproject/demo_share@snap1>
```

2. Enter the `list clones` command.

   ```
   hostname:shares myproject/demo_share@snap1> list clones

   Clones: 2 total

   PROJECT          SHARE
   myproject        demo_clone1
   myproject        demo_clone2
   hostname:shares myproject/demo_share@snap1
   ```

   The result shows how many clones exist, the project in which each resides, and the name of each clone.

Viewing a Clone Origin (BUI)

Use the following procedure to determine the snapshot from which a clone was made.

1. Go to the clone.
   a. From the **Shares** menu, select **Shares**.
   b. Hover over the clone, and click the edit icon ✎ .
2. Under **Static Properties** on the left, click **Show** next to **Clone origin**.

   A window appears giving the name of the snapshot from which the clone was made.

Viewing a Clone Origin (CLI)

Use the following procedure to determine the snapshot from which a clone was made.

1. Go to `shares` and select the project that contains the clone, then select the clone.

   ```
   hostname:> shares select myproject
   hostname:shares myproject> select demo_clone
   hostname:shares myproject/demo_clone>
   ```

2. Enter the `get origin` command.

   The command returns the location and name of the snapshot from which the clone was made.

   ```
   hostname:shares myproject/demo_clone> get origin
   origin = myproject/demo_share@demo_snapshot
   ```

7

Remote Replication

> **Note:**
>
> Replication and Cloning are licensed features for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Oracle ZFS Storage Appliance supports snapshot-based replication of projects and shares from a source appliance to a replication target, to a different pool on the same appliance, or to an NFS server for offline replication. You can configure replication to be executed manually, on a schedule, or continuously. Use cases for remote replication include disaster recovery, data distribution, disk-to-disk backup, and data migration between appliances when upgrading hardware or rebalancing storage.

To configure, monitor, and manage remote replication, use the following tasks:

- Remote Replication Workflow
- Configuring Remote Replication
- Monitoring Remote Replication
- Managing Replication Packages
- Disaster Recovery with Remote Replication

For details about remote replication, see Remote Replication Concepts.

Remote Replication Workflow

The following steps summarize the basic steps for using remote replication. For information about remote replication concepts, see Remote Replication Concepts.

1. Check software compatibility on source and target appliances.

   For information about software compatibility, see Checking Source and Target Compatibility.

2. Set up network interfaces and routing.

   For information about setting up network routing, see Setting Up Network Interfaces and Static Routing - BUI, CLI.

3. Create a replication target.

   For information about creating a replication target, see Creating a Replication Target - BUI, CLI.

4. Create a replication action.

   For information about creating a replication action, see Creating a Replication Action - BUI, CLI.

5. Send a replication update as specified by the replication action.

   Replication updates occur at a specified frequency, on a continuous basis, or manually. See Replication Update Frequency.

6. Optionally, configure offline replication.

   For information about offline replication, see Configuring Offline Replication - BUI, CLI.

Configuring Remote Replication

Use the following tasks to configure remote replication:

- Checking Source and Target Compatibility
- Setting Up Network Interfaces and Static Routing - BUI, CLI
- Creating a Replication Target - BUI, CLI
- Creating a Replication Action - BUI, CLI
- Configuring Automatic Snapshot Management on Target - BUI, CLI
- Manually Sending a Replication Update - BUI, CLI
- Configuring Replication for a Clustered Configuration
- Configuring Offline Replication - BUI, CLI
- Disabling Replication Compression - BUI, CLI
- Testing the Connection - BUI, CLI
- Editing a Replication Target - BUI, CLI
- Editing a Replication Action - BUI, CLI

Checking Source and Target Compatibility

Remote replication is compatible between most Oracle ZFS Storage Appliance software versions. Compatibility failures are caused if a replication update uses a feature that is not supported on the replication target. Features are delivered with software updates or as deferred updates.

For details about compatibility and deferred update features for each software version, see the Oracle ZFS Storage Appliance Remote Replication Compatibility document (Doc ID 1958039.1) on My Oracle Support (https://support.oracle.com/).

1. Check the current software version on the source and target appliances.
   - **BUI** - From the **Maintenance** menu, select **System**.
   - **CLI** - Navigate to `maintenance system updates` and enter `show`.

2. Ensure that the replication target provides support for any deferred update feature used by the source project or share.

   For example, if the source share uses large blocks, ensure that the replication target provides support for this feature. As another example, if the source share uses the mandatory file retention policy, the replication target pool must have a redundant profile and, therefore, the striped profile cannot be used.

3. Update software and apply deferred updates on the replication target, as needed.

   For more information, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

4. If the project or share to be replicated is encrypted, ensure that the encryption key used at the source is also available at the target.

   For information about configuring encryption keys, see Data Encryption.

5. Optional: On the source and target appliances, set the SSL/TLS versions and ciphers for replication as described in Configuring SSL/TLS Versions and Ciphers.

   Set the values according to your site's security requirements. The versions and at least one of the ciphers must be identical on the source and target appliances. Oracle ZFS Storage Appliance systems running older firmware might not support ciphers offered in newer TLS versions.

Setting Up Network Interfaces and Static Routing (BUI)

To ensure the appropriate network interfaces are used for the replication connections between source and target appliances, configure IPv4 static /32 or IPv6 static /128 (host-specific) routes.

If you are setting up replication for a cluster configuration, select a singleton (unlocked) network interface so that following a cluster takeover or failback, the interface will move to the node where the replication work is being done.

1. Source appliance: From the **Configuration** menu, select **Network**, then **Routing**.
2. Click the add icon ⊕ next to **Routing Table Entries**.
3. In the **Insert Static Route** dialog box, define the static route to the target appliance:
   - **Family** - Select **IPv4** or **IPv6**.
   - **Kind** - Select **Network**.
   - **Destination** - For the target appliance, enter the IPv4 address and netmask /32 or the IPv6 address and netmask /128.
   - **Gateway** - For the target appliance, enter the gateway address.
   - **Interface** - Select the interface name from the pull-down menu.

   IPv4 Configuration Example

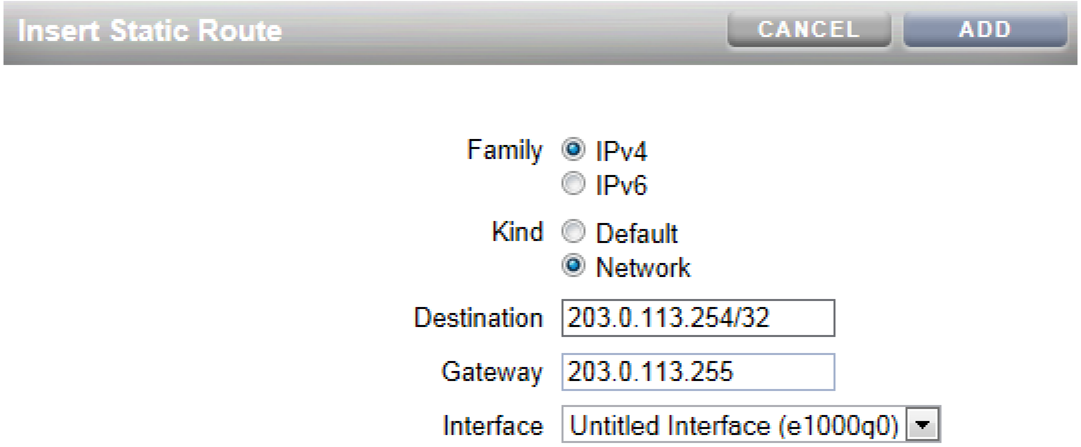

Family	● IPv4 ○ IPv6
Kind	○ Default ● Network
Destination	203.0.113.254/32
Gateway	203.0.113.255
Interface	Untitled Interface (e1000g0) ▾

   IPv6 Configuration Example

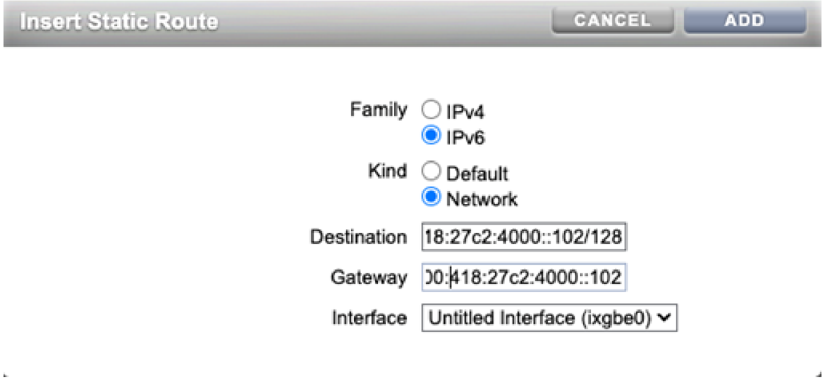

Family ○ IPv4 ● IPv6
Kind ○ Default ● Network
Destination 18:27c2:4000::102/128
Gateway 0:418:27c2:4000::102
Interface Untitled Interface (ixgbe0)

4. Click **ADD**.
5. After defining the static route from the source appliance to the target appliance, repeat these steps on the target appliance to define the static route from the target back to the source.
6. To verify that traffic is routed through the correct source and target interfaces, use the `traceroute` command.

   For information about using `traceroute`, see Configuring Network Routing.

   > **Note:**
   >
   > When an interface is deleted, all routes associated with the interface are also removed.

Related Topics

- Example: Replication Configuration for Clustered Appliances
- Remote Replication Workflow
- Remote Replication Concepts

Setting Up Network Interfaces and Static Routing (CLI)

To ensure the appropriate network interfaces are used for the replication connections between source and target appliances, configure IPv4 static /32 or IPv6 static /128 (host-specific) routes.

If you are setting up replication for a cluster configuration, select a singleton (unlocked) network interface so that following a cluster takeover or failback, the interface will move to the node where the replication work is being done.

1. Source appliance: Go to `configuration services routing`. Define the static route to the target appliance using the appropriate example:

   IPv4 Configuration Example

   Use a static /32 (host-specific) route to the target system IPv4 address via the dedicated network interface. In the following example, `mask=32` means that this is a host-specific route.

   ```
   host_source:configuration services routing> create
   ```

```
    host_source:configuration services route (uncommitted)> get
            family = (unset)
       destination = (unset)
              mask = (unset)
           gateway = (unset)
         interface = (unset)
  host_source:configuration services route (uncommitted)> set family=IPv4
  host_source:configuration services route (uncommitted)> set 
destination=203.34.56.78
  host_source:configuration services route (uncommitted)> set mask=32
  host_source:configuration services route (uncommitted)> set 
gateway=203.34.56.254
  host_source:configuration services route (uncommitted)> set interface=nge3
  host_source:configuration services route (uncommitted)> commit
  host_source:configuration services routing> show
  route-000  0.0.0.0/0         203.24.30.254  nge0  static
  route-001  203.24.30.0/32    203.24.30.28   nge0  dynamic
  route-002  203.24.150.0/32   203.24.150.10  ibd0  dynamic
  route-003  203.24.101.65/32  203.24.30.254  nge1  inactive
  route-005  203.34.56.78/32   203.34.56.254  nge3  static
```

IPv6 Configuration Example

Use a static /128 (host-specific) route to the target system IPv6 address via the dedicated network interface. In the following example, `mask=128` means that this is a host-specific route.

```
host_source:configuration services routing> create

    host_source:configuration services route (uncommitted)> get
            family = (unset)
       destination = (unset)
              mask = (unset)
           gateway = (unset)
         interface = (unset)
  host_source:configuration services route (uncommitted)> set family=IPv6
  host_source:configuration services route (uncommitted)> set 
destination=2606:b400:418:27c2:4000::102
  host_source:configuration services route (uncommitted)> set mask=128
  host_source:configuration services route (uncommitted)> set 
gateway=2606:b400:418:27c2:4000::102
  host_source:configuration services route (uncommitted)> set interface=ixgbe0
  host_source:configuration services route (uncommitted)> commit
  host_source:configuration services routing> show
  route-000 0.0.0.0/0         203.24.30.254  nge0   static
  route-001 203.24.30.0/32    203.24.30.28   nge0   dynamic
  route-002 203.24.150.0/32   203.24.150.10  ibd0   dynamic
  route-003 203.24.101.65/32  203.24.30.254  nge1   inactive
  route-005 203.34.56.78/32   203.34.56.254  nge3   static
  route-006 2606:b400:418:27c2:4000::102/128 2606:b400:418:27c2:4000::102 ixgbe0 
static active
```

2. After defining the static route from the source appliance to the target appliance, repeat these steps on the target appliance to define the static route from the target back to the source.

3. To ensure traffic is routed through the correct source and target interfaces, use the `traceroute` command.

   For information about using `traceroute`, see Configuring Network Routing.

Note:

When an interface is deleted, all routes associated with the interface are also removed.

Related Topics

- Example: Replication Configuration for Clustered Appliances
- Remote Replication Workflow
- Remote Replication Concepts

Creating a Replication Target (BUI)

A replication target establishes a secure communication connection between source and target appliances.

Before You Begin

See Checking Source and Target Compatibility to ensure your replication target is compatible with the source.

If you need to ensure that the replication traffic goes over a particular network interface, set up a static route for the target that specifies that interface as shown in Setting Up Network Interfaces and Static Routing (BUI).

1. Source appliance: From the **Configuration** menu, select **Services**, then **Remote Replication**.
2. Next to **Targets**, click the add icon ⊕ .

   The **Add Replication Target** dialog box is displayed.
3. Enter information about the replication target.
   - **Name** - The name of the target to display in the BUI and CLI of the source appliance.
   - **Hostname** - The fully qualified domain name, or IPv4 or IPv6 address of the target appliance. The recommended value to use is the target's fully qualified domain name. See the description for **Host Match**.
   - **Username** - The name of the user on the target appliance who is authorized to set up replication relations (**Appliance**/**peerSetup** authorization). It is not recommended to use the "root" user account.
   - Choose one of the following authentication options and complete the field:

     Password - The password for the user specified in the **Username** field.

     Token - The REST authorization token for the user specified in the **Username** field. For information about REST tokens, see Preference Properties.
   - **Host Match** - When this option is enabled, the system verifies that the hostname that you specified in the **Hostname** field in this dialog box matches a host specified in the certificate. If you specify an IP address or an unqualified domain name for the **Hostname**, and the certificate only has fully qualified domain names, the **Host Match** fails and the target is not created.

     If the target is using an ASN-based certificate, specify the target's fully qualified domain name for the value of the **Hostname** property.

     If you disable the **Host Match** option, hostname validation is not performed.

For stronger security, set the value of **Hostname** to the target's fully qualified domain name, and make sure the **Host Match** option is enabled.

4. Click **Add**.

   The certificate trust check is performed to verify whether the certificate is trusted.

   If the certificate is not trusted by the source, the certificate is presented for you to review, and you are prompted to accept or reject the certificate as described in Testing the Connection (BUI). If you accept the certificate, the certificate is added to the trust list of the source, and the target is created. If you reject the certificate, the certificate is not added to the trust list of the source, and the target is not created.

   If the certificate is already trusted, the target is created, and you are not prompted to accept the certificate.

Related Topics

- Testing the Connection (BUI)
- Editing a Replication Target (BUI)
- Replication Targets
- Remote Replication Workflow
- Remote Replication Concepts
- Backing Up, Replicating, and Restoring Encrypted Projects and Shares

Creating a Replication Target (CLI)

A replication target establishes a secure communication connection between source and target appliances.

Before You Begin

See Checking Source and Target Compatibility to ensure your replication target is compatible with the source.

If you need to ensure that the replication traffic goes over a particular network interface, set up a static route for the target that specifies that interface as shown in Setting Up Network Interfaces and Static Routing (CLI).

1. Source appliance: Go to `configuration services replication targets`.
2. Enter the `target` command.

   ```
   host_source:> configuration services replication targets> target
   host_source:configuration services replication target (uncommitted)>
   ```

3. Set the target properties.

   - `hostname` - The fully qualified domain name, or IPv4 or IPv6 address of the target appliance. The recommended value to use is the target's fully qualified domain name. See the description for `host_match`.
   - `user` - The name of the user on the target appliance who is authorized to set up replication relations (`appliance`/`allow_peerSetup` authorization). It is not recommended to use the "root" user account.
   - Choose one of the following authentication options and set its value:

     `password` - The password for the user specified in the `user` property.

token - The REST authorization token for the user specified in the user property. For information about REST tokens, see Preference Properties.

- label - The name of the target to display in the BUI and CLI of the source appliance.
- host_match - When this property is true, the system verifies that the target hostname specified in the hostname property matches the host specified in the certificate. For example, if the certificate subject common name only has a domain name, and if you specify an IP address for hostname, this hostname check fails. If the hostname check fails, the certificate trust check described in the following step is not performed and the target is not created.

  If the target is using an ASN-based certificate, specify the target's fully qualified domain name for the value of the hostname property.

  If you set host_match to false, hostname validation is not performed.

  For stronger security, set the value of the hostname property to the target's fully qualified domain name, and make sure the host_match property is set to true.

```
hostname:configuration services replication target (uncommitted)> set
hostname=hostname
hostname:configuration services replication target (uncommitted)> set
root_password=pw
hostname:configuration services replication target (uncommitted)> set label=repl_1
hostname:configuration services replication target (uncommitted)> set
host_match=true
```

4. Commit the changes.

   The certificate trust check is performed to verify whether the certificate is trusted.

   If the certificate is not trusted by the source, the certificate is presented for you to review, and you are prompted to accept or reject the certificate as described in Testing the Connection (CLI). If you accept the certificate, the certificate is added to the trust list of the source, and the target is created. If you reject the certificate, the certificate is not added to the trust list of the source, and the target is not created.

   If the certificate is already trusted, the target is created, and you are not prompted to accept the certificate.

Related Topics

- Testing the Connection (CLI)
- Editing a Replication Target (CLI)
- Replication Targets
- Remote Replication Workflow
- Remote Replication Concepts
- Backing Up, Replicating, and Restoring Encrypted Projects and Shares

Creating a Replication Action (BUI)

A replication action describes the project or share to be replicated, where to send the replication, the replication schedule, and data transfer properties such as enabling or disabling encryption of the network link.

Caution:

Do not create more than three actions per project. One action per project is typical.

If you are setting up remote replication for the first time, it might be useful to replicate a minimal amount of data to ensure the synchronization completes successfully. You can either replicate an empty project or choose not to replicate the snapshots in the project/shares.

If you are replicating a large data set and bandwidth is limited due to distance between source and target appliances, you can export the replication to offline media, as described in Configuring Offline Replication (BUI).

1. Source appliance: From the **Shares** menu, select **Projects**.
2. Select the project or share, and click the **Replication** tab.
3. Next to **Actions**, click the add icon ⊕.
4. Select a target and a pool.
5. Select properties for this action.

   See Replication Action Properties for a description of all properties.
6. Select **Scheduled**, and set a frequency for the replication update, or select **Continuous** to send replication updates continuously.
7. Click **Add**.

   The replication action is added to the **Actions** list.

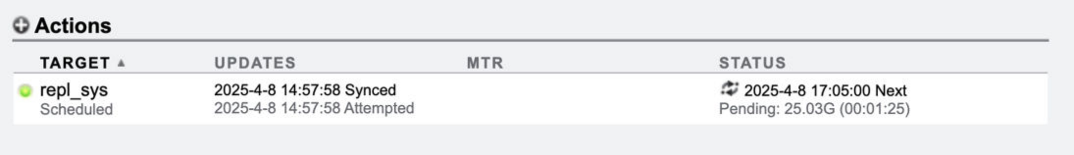

⊕ Actions

TARGET ▲	UPDATES	MTR	STATUS
● repl_sys Scheduled	2025-4-8 14:57:58 Synced 2025-4-8 14:57:58 Attempted		2025-4-8 17:05:00 Next Pending: 25.03G (00:01:25)

   The Status column of the **Actions** list shows the scheduled time for the next replication update. The second line shows estimates for the data size and the data transfer time for the replication update. For example, the graphic shows that the estimated data size of the replication update is `25.03G` and the estimated data transfer time is `00:01:25`.

Related Topics

- Replication Action Properties
- Replication Actions and Packages
- Manually Sending a Replication Update (BUI)

Creating a Replication Action (CLI)

A replication action describes the project or share to be replicated, the replication target, the replication schedule, and data transfer properties such as enabling or disabling encryption of the network link.

Caution:

Do not create more than three actions per project. One action per project is typical.

If you are setting up remote replication for the first time, it might be useful to replicate a minimal amount of data to ensure the sync completes successfully. You can either replicate an empty project or choose not to replicate the snapshots in the project/shares.

If you are replicating a large data set and bandwidth is limited due to distance between source and target appliances, you can export the replication to offline media, as described in Configuring Offline Replication (CLI).

1. Source appliance: Navigate to the project or share, and enter command `action`.

   ```
   host_source:shares PROJECT1 replication> action
   ```

2. Display the properties using the `get` command.

   ```
   host_source:shares PROJECT1 action (uncommitted)> get
                           target = (unset)
                             pool = (unset)
                          enabled = true
                       continuous = false
              update_cascade_delay =
                    include_snaps = true
       retain_user_snaps_on_target = false
                            dedup = false
       include_clone_origin_as_data = false
                    max_bandwidth = unlimited
                       bytes_sent = 0
                   estimated_size = 0
              estimated_time_left = 00:00:00
      pending_update_estimated_size = 0
      pending_update_estimated_time = 00:00:00
               average_throughput = 0B/s
                          use_ssl = true
                      compression = on
                      export_path =
         recovery_point_objective =
          replica_lag_warning_alert =
            replica_lag_error_alert =
                 potential_source = false
                   distant_target = false
                        rawcrypto = off
   ```

3. Set the properties for this action.

   See Replication Action Properties for a description of CLI properties.

   ```
   host_source:shares PROJECT1 action-000 (uncommitted)> set target=repl_sys
                            target = repl_sys (uncommitted)
   host_source:shares PROJECT1 action-000 (uncommitted)> set pool=pool-0
                                pool = pool-0 (uncommitted)
   host_source:shares PROJECT1 action-000 (uncommitted)> set include_snaps=false
                        include_snaps = false (uncommitted)
   host_source:shares PROJECT1 action-000 (uncommitted)> set use_ssl=false
                          use_ssl = false (uncommitted)
   host_source:shares PROJECT1 action-000> schedule
   host_source:shares PROJECT1 action-000 schedule (uncommitted)> set frequency=day
                               frequency = day (uncommitted)
   host_source:shares PROJECT1 action-000 schedule (uncommitted)> set hour=23
   ```

```
                                 hour = 23 (uncommitted)
host_source:shares PROJECT1 action-000 schedule (uncommitted)> set minute=05
                               minute = 05 (uncommitted)
```

4. Commit the new replication action.

```
host_source:shares PROJECT1 action-000 schedule (uncommitted)> commit
```

5. To view the properties of the newly created action, enter `ls`:

```
host_source:shares PROJECT1 replication> ls
 Properties:
      inherited = false
 Actions:
      TARGET                          STATUS     NEXT
     action-000   repl_sys            idle       manual

host_source:shares PROJECT1 action-000> ls
Properties:

                            id = 49b1ec85-cb8b-436f-8a56-87189c18e066
                     target_id = 45f8a28e-218d-4634-94a2-c913cc582a44
                        target = repl_sys
                  target_pkgid = 49b1ec85-cb8b-436f-8a56-87189c18e066
                   target_pool = pool-0
                       enabled = true
                    continuous = false
          update_cascade_delay =
                 include_snaps = false
   retain_user_snaps_on_target = false
                         dedup = false
  include_clone_origin_as_data = false
                 max_bandwidth = unlimited
                    bytes_sent = 0
                estimated_size = 0
           estimated_time_left = 00:00:00
 pending_update_estimated_size = 25.0329161G
 pending_update_estimated_time = 00:01:25
            average_throughput = 0B/s
                       use_ssl = false
                   compression = on
                   export_path =
                         state = idle
             state_description = Idle (no update in progress)
                export_pending = false
                       offline = false
                   next_update = Tue Apr 08 2025 23:05:00 GMT+0000 (UTC)
        replica_data_timestamp = Tue Apr 08 2025 20:56:51 GMT+0000 (UTC)
                     last_sync = Tue Apr 08 2025 20:57:58 GMT+0000 (UTC)
                      last_try = Tue Apr 08 2025 20:57:58 GMT+0000 (UTC)
                   last_result = success
                   replica_lag = 00:03:33
      recovery_point_objective =
     replica_lag_warning_alert =
       replica_lag_error_alert =
replica_lag_over_warning_limit = false
  replica_lag_over_error_limit = false
              potential_source = false
                distant_target = false
                     rawcrypto = off

Schedules:
```

```
NAME                 FREQUENCY            DAY                  HH:MM
schedule-000         day                  -                    23:05

Children:
                        autosnaps => Configure automatic snapshots on target
                        schedules => Configure replication update schedules
                            stats => Replication Action Statistics Properties
```

The `pending_update_estimated_size` property shows the estimated data size for the pending replication update (25.0329161G). The `pending_update_estimated_time` property shows the estimated data transfer time for the pending replication update (00:01:25).

6. To view the ID of the newly created action, use the `last` command, which navigates to the node with the new action, combined with `get id`, which retrieves the action ID.

   The ID is used later to select the correct replication action node.

   ```
   host_source:shares PROJECT1 replication> last get id
                                  id = 49b1ec85-cb8b-436f-8a56-87189c18e066
   ```

Related Topics

- Editing a Replication Target (CLI)
- Replication Action Properties
- Replication Actions and Packages

Configuring Automatic Snapshot Retention on a Target (BUI)

Use this procedure to set a different number of retained automatic snapshots on the replication target than what was set on the source appliance. Before performing this procedure, you must set a replication action on a project or share, and schedule automatic snapshots.

When performing this task, modify the replication action and the snapshot schedule from the appropriate project or share. If the replication action and snapshot schedule are set at different levels, edit the replication action at the same level where the schedule is configured, as shown in the following table:

Action	Snapshot Schedule	Modify automatic snapshot retention on:
Project level	Project level	Action configured at the project level.
Share level	Share level	Action configured at the share level.
Project level	Share level	Action visible at the share level. (Project-level action inherited by share.)
Share level	Project level	Action configured at the share level.

See Replication Automatic Snapshot Management for more information.

Before You Begin

Before configuring automatic snapshot retention on a target, you must first do the following:

- Create a replication action on a project or share: Creating a Replication Action (BUI).
- Create an automatic snapshot schedule for the project or share: Scheduling Snapshots (BUI).

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to make changes to a retention policy:

- Modify an automatic snapshot schedule: `scheduleSnap`
- Modify the retention hold value from **Off** to **Locked**: `scheduleLockedSnap`
- Modify the retention hold value from **Locked** to **Off**: `releaseSnapRetention`

1. Source appliance: From the **Shares** menu, select **Projects**.
2. Select the project or share with the replication action and snapshot schedule, and click the **Replication** tab.
3. Click the edit icon ✎ of the action you want to modify.
4. At the bottom of the **Edit Replication Action** window, click the **Snapshots** tab.

   The automatic snapshot schedules appear.
5. For field **Auto Snapshot Retention Policies**, select **independent**.
6. For field **Keep At Most**, set the number of automatic snapshots to be retained on the target.
7. For field **Retention**, select a retention policy option:
   - **Off -** No retention policy applies to this snapshot.
   - **Locked** - To lock the snapshot schedule, select **Locked** for the retention policy, and set **Retention** to the number of snapshots that should be locked and, thus, protected from manual deletion. The number for the **Retention** property must be the same or smaller than the number for the **Keep At Most** property. When all locked snapshots for this schedule have exceeded the **Keep At Most** property value, the retention hold changes from **Locked** to **Off**.
8. Click **Apply**.

> **Note:**
>
> Automatic snapshot retention for replication has special processing during reverse replication. For more information, see Replication Snapshot Management.

Configuring Automatic Snapshot Retention on a Target (CLI)

Use this procedure to set a different number of retained automatic snapshots on the replication target than what was set on the source appliance. Before performing this procedure, you must set a replication action on a project or share, and schedule automatic snapshots.

When performing this task, modify the replication action and the snapshot schedule from the appropriate project or share. If the replication action and snapshot schedule are set at different levels, edit the replication action at the same level where the schedule is configured, as shown in the following table:

Action	Snapshot Schedule	Modify automatic snapshot retention on:
Project level	Project level	Action configured at the project level.
Share level	Share level	Action configured at the share level.
Project level	Share level	Action visible at the share level. (Project-level action inherited by share.)
Share level	Project level	Action configured at the share level.

See Replication Automatic Snapshot Management for more information.

Before You Begin

Before configuring automatic snapshot retention on a target, you must first do the following:

- Create a replication action on a project or share: Creating a Replication Action (CLI).
- Create an automatic snapshot schedule for the project or share: Scheduling Snapshots (CLI).

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to make changes to a retention policy:

- Modify an automatic snapshot schedule: `scheduleSnap`
- Modify the retention hold value from `off` to `locked`: `scheduleLockedSnap`
- Modify the retention hold value from `locked` to `off`: `releaseSnapRetention`

1. Source appliance: Go to `shares` and select the appropriate project or share with the replication action and snapshot schedule.

   ```
   hostname:> shares select MyProject
   hostname:shares MyProject> select MyShare
   hostname:shares MyProject/MyShare>
   ```

2. Enter `replication`, then enter `show` to display existing actions.

   ```
   hostname:shares MyProject/MyShare> replication
   hostname:shares MyProject/MyShare replication> show
   Properties:
                         inherited = false

   Actions:

               TARGET          STATUS     NEXT
   action-000  local           idle       Thu Jan 03 2019 12:04:00 GMT+0000 (UTC)
   ```

3. Select the action for which you want to modify retention settings.

   ```
   hostname:shares MyProject/MyShare replication> select action-000
   hostname:shares MyProject/MyShare action-000>
   ```

4. Enter `autosnaps`, then enter `show` to display snapshot schedules.

   Note if property `autosnaps_retention_policies` is set to `synchronized` or `independent`.

   ```
   hostname:shares MyProject/MyShare action-000> autosnaps
   hostname:shares MyProject/MyShare action-000 autosnaps> show
   Properties:
   ```

```
  autosnaps_retention_policies = independent

Automatics:

NAME            FREQUENCY       DAY              HH:MM KEEP
automatic-000   hour            -                 -:04    3
```

5. If property `autosnaps_retention_policies` is set to `synchronized`, change it to `independent` and enter `commit`.

```
hostname:shares MyProject/MyShare action-000 autosnaps> set
autosnaps_retention_policies=independent
                           autosnaps_retention_policies = independent (uncommitted)
hostname:shares MyProject/MyShare action-000 autosnaps> commit
```

6. Select the automatic snapshot schedule that you want to modify.

```
hostname:shares MyProject/MyShare action-000 autosnaps> select automatic-000
hostname:shares MyProject/MyShare action-000 automatic-000>
```

7. Use the `ls` command to show the schedule's properties. Note if property `retentionpolicy` is set to `off` or `locked`.

```
hostname:shares MyProject/MyShare action-000 automatic-000> ls
Properties:
                      frequency = hour
                            day = 
                           hour = 
                         minute = 04
                           keep = 10
                  retentionhold = 0
                retentionpolicy = off
```

8. Set property `keep` to the number of automatic snapshots to be retained on the target.

```
hostname:shares MyProject/MyShare action-000 automatic-000> set keep=10
                           keep = 10 (uncommitted)
```

9. Set the retention policy property to one of the following options:

   - `off` - To not set a retention policy for this snapshot, set property `retentionpolicy` to `off`.
   - `locked` - To lock the snapshot schedule, set property `retentionpolicy` to `locked`, and set property `retentionhold` to the number of snapshots that should be locked and, thus, protected from manual deletion. The number for the `retentionhold` property must be the same or smaller than the number for the `keep` property. When all locked snapshots for this schedule have exceeded the `keep` property value, the retention hold changes from `locked` to `off`.

```
hostname:shares MyProject/MyShare action-000 automatic-000> set
retentionpolicy=locked
                           retentionpolicy = locked (uncommitted)
hostname:shares MyProject/MyShare action-000 automatic-000> set retentionhold=2
                           retentionhold = 2 (uncommitted)
```

10. Enter `commit` to save the latest changes, and enter `show` to verify the changes.

```
hostname:shares MyProject/MyShare action-000 automatic-000> commit
hostname:shares MyProject/MyShare action-000 automatic-000> show
Properties:
                      frequency = hour
                            day = 
                           hour = 
                        minute = 04
                          keep = 10
                 retentionhold = 2
               retentionpolicy = locked
```

> **Note:**
>
> Automatic snapshot retention for replication has special processing during reverse replication. For more information, see Replication Snapshot Management.

Manually Sending a Replication Update (BUI)

If continuous or scheduled replication is already configured, replication updates are performed automatically. You can also perform a manual update using the BUI.

1. Source appliance: From the **Shares** menu, select **Projects**.
2. Open a project, and click the **Replication** tab.
3. Click the Sync Now icon ⟳ .

> **Note:**
>
> This action is not available (or will not work) if an update is actively being sent. Ensure there is enough disk space on the target to replicate the entire project before sending an update.

The BUI displays a progress bar and indicates when the update completes. If the replication update does not complete successfully, click the Sync Now icon again. The replication resumes from the point of interruption.

Related Topics

- Creating a Replication Action (BUI)
- Canceling a Replication Update (BUI)

Manually Sending a Replication Update (CLI)

If continuous or scheduled replication is already configured, replication updates are performed automatically. You can also perform a manual update using the CLI.

1. Source appliance: Go to `shares`.
2. Select the share.
3. Enter the `sendupdate` command.

```
host_source:shares PROJECT1/SHARE1 action-000> sendupdate
```

When the update is currently active, the CLI shows a state of `sending`.

If the replication update does not complete successfully, enter the `sendupdate` command again. The replication resumes from the point of interruption.

Related Topics

- Creating a Replication Action (CLI)
- Canceling a Replication Update (CLI)

Configuring Replication for a Clustered Configuration

This task describes how to configure replication in a clustered environment. Follow these steps to configure replication properly to ensure that projects continue to replicate after a cluster takeover, cluster failback, or after performing reverse replication on a target appliance.

Before You Begin

If you are configuring replication for clustered appliances for the first time, review Example: Replication Configuration for Clustered Appliances.

1. On the replication source and target appliances, select network interfaces and IP addresses to be used for replication traffic, using these guidelines:
   a. Always select a singleton network interface to ensure it is taken over by the peer node after a cluster takeover or failback operation.
   b. On the source system, ensure that the selected network interface and the storage pool, from which the data will be replicated, are both assigned to the same node. This is always the case when the source cluster is in the `CLUSTERED` state.
   c. On the target system, assign the selected network interface on the target appliance and the storage pool, into which the replicated data will be put, to the same node. This maintains the association when the replication configuration is performed while the target cluster is in the `CLUSTERED` state.
   d. Ensure that the source and the target systems can communicate using the selected network interfaces and IP addresses.
2. On the source and target appliances, create static /32 (host-based) network routes using the selected network interfaces and IP addresses.
3. On the source appliance, configure the replication target object using the selected IP address of the target.

Configuring Offline Replication (BUI)

Use the following steps to configure offline replication:

- Setting Up an NFS Server for Offline Replication
- Setting Up an Export Path to the NFS Server (BUI)
- Exporting a Replication Update (BUI)
- Verifying the Replication Stream On the NFS Server
- Importing the Replication Stream from the NFS Server (BUI)
- Performing a Manual Network Update (BUI)
- Reversing an Offline Replication (BUI)

Setting Up an NFS Server for Offline Replication

The steps for setting up an NFS server will vary depending on the NFS server type you use. Refer to your NFS server documentation for specific instructions.

1. Identify a server that is network ready and has NFS services enabled.

2. As root of the NFS server, create a filesystem or share.
3. Set the file permissions to expose the NFS share only to the IP address of the source and target appliances.
4. To encrypt the replication stream, enable on-disk encryption for the NFS share on the NFS server.

> **Note:**
>
> An exported replication stream is never encrypted by the appliance.

5. Export the share for access by the NFS client.
6. Verify that the filesystem is shared.

Next Steps

Setting Up an Export Path to the NFS Server (BUI)

Setting Up an Export Path to the NFS Server (BUI)

Use the following procedure to set up an export path to the NFS server.

Before You Begin

Identify or create a target. To create a target, see Creating a Replication Target (BUI).

1. Source appliance: From the **Shares** menu, select **Projects**.
2. Open the project, and click the **Replication** tab.
3. Click the add icon ⊕ next to **Actions**.
4. In the **Add Replication Action** dialog box, select **Export data path** and enter the path of the NFS share in the form `nfs://server/path`.
5. Select additional properties for this action, and then click **Add**.

> **Note:**
>
> If you configure a schedule or select continuous replication mode, the update will occur automatically after the export and the import operations have completed.

Next Steps

Exporting a Replication Update (BUI)

Exporting a Replication Update (BUI)

Use the following procedure to export a replication update.

1. Source appliance: From the **Shares** menu, select **Projects**.
2. Open the project, and click the **Replication** tab.
3. Click the Export Replication Data icon .
4. Check the replication status, and wait until the replication completes.

Next Steps

Verifying the Replication Stream On the NFS Server

Verifying the Replication Stream on the NFS Server

Use the following procedure to verify the replication stream on the NFS server.

1. Navigate to the NFS directory, check the MD5 stream, and view the metadata.

```
hostname# pwd
/export/init_repl/rr_updates/96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
hostname# ls -l
total 67
-rw-r--r--   1 nobody   nobody       633 Nov 17 21:46 metadata.xml
-rw-------   1 nobody   nobody     31016 Nov 17 21:46 stream
-rw-------   1 nobody   nobody        33 Nov 17 21:46 stream.md5
hostname# md5sum stream
25b4671c9aaf34455a63e203bcecff49  stream
hostname# cat stream.md5
25b4671c9aaf34455a63e203bcecff49
hostname# cat metadata.xml
<?xml version="1.0"?>
<!DOCTYPE nvlist SYSTEM "/usr/share/lib/xml/dtd/nvlist.dtd.1">
<nvlist>
   <nvpair name='offline_rr_version'><string value='1.1'/></nvpair>
   <nvpair name='source_asn'><string value='2ea4670f-bc17-cf8f-a420-9211d6edda04'/></
nvpair>
   <nvpair name='project'><string value='default'/></nvpair>
   <nvpair name='pkgid'><string value='96366bf2-0b3c-4eec-e85b-e36e1b5bc18c'/></
nvpair>
   <nvpair name='basesnap'><string value=/></nvpair>
   <nvpair name='newsnap'><string value='.rr-96366bf2-0b3c-4eec-e85b-
e36e1b5bc18c-1'/></nvpair>
   <nvpair name='compression'><string value='on'/></nvpair>
</nvlist>
hostname#
```

2. Physically move the NFS server to the target appliance site, or copy the `rr_updates` folder to external media and prepare for shipping.

Next Steps

Importing the Replication Stream from the NFS Server (BUI)

Importing the Replication Stream from the NFS Server (BUI)

Use the following procedure to import the replication stream from the NFS server.

1. From the **Shares** menu, select **Projects**, then **Replicas**.
2. Select the replica that shows the source is `awaiting import`.
3. Click the **Replication** tab.
4. In the **Import Data Path** field, enter the path of the replica.
5. Click the Import Update from External Media icon ⤒ to start the import.

Next Steps

After the replication stream is imported to the target appliance, continue with one of the following procedures:

- Performing a Manual Network Update (BUI)

- Reversing an Offline Replication (BUI)

Performing a Manual Network Update (BUI)

After importing the offline replication stream to the target appliance, confirm future network updates will work correctly. If continuous or scheduled replication is already configured, the update will be performed automatically. Otherwise, perform a manual update.

1. Go to the source appliance.
2. See Manually Sending a Replication Update (BUI).

Reversing an Offline Replication (BUI)

Follow this procedure to move an offline replication package to a new local project, configured to replicate back to a source appliance.

1. Import the offline replication package from an NFS server to the target appliance, as described in Importing the Replication Stream from the NFS Server (BUI).
2. Target appliance: From the **Shares** menu, select **Projects**, then **Replica**, and locate the replicated package.

   The project is named *target_appliance*: `new_project/share`.
3. Select the project, and click its Reverse Replication Direction icon .
4. In the **Reverse Replication** window, enter a name for the new local project.

   This action moves the contents of this package to a new local project configured to replicate back to the source. Any data or metadata changes made on the source since the last successful update will be lost when the new project is replicated back to the source. If replication actions on the source are not disabled, future updates to this package will fail.
5. From the **Shares** menu, select **Projects**.
6. Open the project, and click the **Replication** tab.
7. Click the Export Replication Data icon .
8. Check the replication status, and wait until the replication completes.
9. After the replication update is complete, navigate to the newly reversed package on the new target.

   The state description should be `Idle (awaiting import)`.
10. Import the update from the NFS server.

Related Topics

Configuring Offline Replication (BUI)

Configuring Offline Replication (CLI)

Use the following steps to configure offline replication:

- Setting Up an NFS Server for Offline Replication
- Setting Up an Export Path to the NFS Server (CLI)
- Exporting a Replication Update (CLI)
- Verifying a Replication Stream On the NFS Server

- Importing a Replication Stream from the NFS Server (CLI)
- Performing a Manual Network Update (CLI)
- Reversing an Offline Replication (CLI)

Setting Up an NFS Server for Offline Replication

The steps for setting up an NFS server will vary depending on the NFS server type you use. Refer to your NFS server documentation for specific instructions.

1. Identify a server that is network ready and has NFS services enabled.
2. As root of the NFS server, create a filesystem or share.
3. Set the file permissions to expose the NFS share only to the IP address of the source and target appliances.
4. To encrypt the replication stream, enable on-disk encryption for the NFS share on the NFS server.

> **Note:**
>
> An exported replication stream is never encrypted by the appliance.

5. Export the share for access by the NFS client.
6. Verify that the filesystem is shared.

Next Steps

Setting Up an Export Path to the NFS Server (CLI)

Setting Up an Export Path to the NFS Server (CLI)

Use the following procedure to set up an export path to the NFS server.

1. Identify or create a replication target.
2. Create a replication action, set the `export_path`, and `commit` the new action.

```
source:shares default replication> action
source:shares default action (uncommitted)> set target=target_a
                        target = target_a (uncommitted)
source:shares default action (uncommitted)> set pool=pool2
                          pool = pool2 (uncommitted)
source:shares default action (uncommitted)> set export_path=nfs://nfs_server/
export/init_repl
                   export_path = nfs://nfs_server/export/init_repl (uncommitted)
source:shares default action (uncommitted)>commit
```

> **Note:**
>
> Optionally, you can set a scheduled or continuous replication mode, which will start the update after the export and the import operations have completed.

3. Navigate back to the replication action that you just created, and view the current status.

```
source:shares default replication> ls
Actions:
               TARGET          STATUS     NEXT
action-000  target_a         idle       Export replication data


source:shares default replication> last
source:shares default action-000> ls
Properties:
                           id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                       target = target_a
                      enabled = true
                   continuous = false
                include_snaps = true
                max_bandwidth = unlimited
                   bytes_sent = 0
               estimated_size = 0
          estimated_time_left = 00:00:00
           average_throughput = 0B/s
                      use_ssl = true
                  compression = on
                  export_path = nfs://nfs_server/export/init_repl
                        state = idle
            state_description = Idle (export pending)
               export_pending = true
                      offline = false
                  next_update = Export replication data
                    last_sync = <unknown>
                     last_try = <unknown>
                  last_result = <unknown>
```

Next Steps

Exporting a Replication Update (CLI)

Exporting a Replication Update (CLI)

Use the following procedure to export a replication update.

1. To export the replication update to the NFS server, use the `sendupdate` command.

   ```
   source:shares default action-000> sendupdate
   ```

2. Enter `ls` to view the status, as shown in this example:

   ```
   source:shares default action-000> ls
   Properties:
                              id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                          target = target_a
                         enabled = true
                      continuous = false
                   include_snaps = true
                   max_bandwidth = unlimited
                      bytes_sent = 0
                  estimated_size = 0
             estimated_time_left = 00:00:00
              average_throughput = 0B/s
                         use_ssl = true
                     compression = on
                     export_path = nfs://nfs_server/export/init_repl
                           state = sending
               state_description = Exporting update
                  export_pending = true
   ```

```
                         offline = false
                     next_update = Export replication data
                       last_sync = <unknown>
                        last_try = <unknown>
                     last_result = <unknown>
```

3. To determine when the export has completed, enter `ls` to view the status.

   Look for `last_result = success`, as shown in this example:

```
source:shares default action-000> ls
Properties:
                              id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                          target = target_a
                         enabled = true
                      continuous = false
                   include_snaps = true
                   max_bandwidth = unlimited
                      bytes_sent = 0
                  estimated_size = 0
             estimated_time_left = 00:00:00
              average_throughput = 0B/s
                         use_ssl = true
                     compression = on
                     export_path =
                           state = idle
               state_description = Idle (no update in progress)
                  export_pending = false
                         offline = true
                     next_update = Sync now
                       last_sync = <unknown>
                        last_try = Mon Nov 18 2019 04:40:40 GMT+0000 (UTC)
                     last_result = success
source:shares default action-000>
```

Next Steps

Verifying a Replication Stream On the NFS Server

Verifying the Replication Stream on the NFS Server

Use the following procedure to verify the replication stream on the NFS server.

1. Navigate to the NFS directory, check the MD5 stream, and view the metadata.

```
hostname# pwd
/export/init_repl/rr_updates/96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
hostname# ls -l
total 67
-rw-r--r--   1 nobody   nobody       633 Nov 17 21:46 metadata.xml
-rw-------   1 nobody   nobody     31016 Nov 17 21:46 stream
-rw-------   1 nobody   nobody        33 Nov 17 21:46 stream.md5
hostname# md5sum stream
25b4671c9aaf34455a63e203bcecff49  stream
hostname# cat stream.md5
25b4671c9aaf34455a63e203bcecff49
hostname# cat metadata.xml
<?xml version="1.0"?>
<!DOCTYPE nvlist SYSTEM "/usr/share/lib/xml/dtd/nvlist.dtd.1">
<nvlist>
   <nvpair name='offline_rr_version'><string value='1.1'/></nvpair>
   <nvpair name='source_asn'><string value='2ea4670f-bc17-cf8f-a420-9211d6edda04'/></
nvpair>
```

```
   <nvpair name='project'><string value='default'/></nvpair>
   <nvpair name='pkgid'><string value='96366bf2-0b3c-4eec-e85b-e36e1b5bc18c'/></
nvpair>
   <nvpair name='basesnap'><string value=/></nvpair>
   <nvpair name='newsnap'><string value='.rr-96366bf2-0b3c-4eec-e85b-
e36e1b5bc18c-1'/></nvpair>
   <nvpair name='compression'><string value='on'/></nvpair>
</nvlist>
hostname#
```

2. Physically move the NFS server to the target appliance site, or copy the `rr_updates` folder to external media and prepare for shipping.

Next Steps

Importing the Replication Stream from the NFS Server (BUI)

Importing a Replication Stream from the NFS Server (CLI)

Use the following procedure to import a replication stream from the NFS server.

1. To import the replication stream from the NFS server, navigate to `shares replication packages` on the target, and then enter `ls` to list the packages.

```
target_a:> shares replication packages
target_a:shares replication packages> ls
Packages:

ID          STATE DATA_TIMESTAMP      SOURCE     DATASET
package-000 idle  unknown             sourceA   <unknown>
```

2. Select the package you want to import. Enter `ls` to view the properties.

```
target_a:shares replication packages> select package-000
target_a:shares replication package-000> ls
Properties:

                            id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                   source_name = sourceA
                    source_asn = d1fce51d-b8a9-6cf8-d71e-fcd4fe42cd0e
                     source_ip = 10.000.000.000:216
                   target_pool = poolA
                    replica_of = <unknown>
                       enabled = true
                         state = idle
             state_description = Idle (no update in progress)
                       offline = false
                   import_path =
                data_timestamp = unknown
                     last_sync = unknown
                      last_try = unknown
                   last_result = unknown
```

3. Set the import path of the replicated data, and then enter `commit`.

```
target_a:shares replication package-000> set import_path=
nfs://nfs_server/export/init_repl
                  import_path = nfs://nfs_server/export/init_repl (uncommitted)
target_a:shares replication package-000> commit
target_a:shares replication package-000> ls
Properties:
```

```
                            id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                   source_name = sourceA
                    source_asn = d1fce51d-b8a9-6cf8-d71e-fcd4fe42cd0e
                     source_ip = 10.000.000.000:216
                   target_pool = poolA
                    replica_of = <unknown>
                       enabled = true
                         state = receiving
             state_description = Importing update
                       offline = true
                   import_path = nfs://nfs_server/export/init_repl
                data_timestamp = unknown
                     last_sync = unknown
                      last_try = unknown
                   last_result = unknown
```

Next Steps

After the replication stream is imported to the target appliance, continue with one of the following procedures:

- Performing a Manual Network Update (CLI)
- Reversing an Offline Replication (CLI)

Performing a Manual Network Update (CLI)

After importing the offline replication stream to the target appliance, confirm future network updates will work correctly. If continuous or scheduled replication is already configured, the update will be performed automatically. Otherwise, perform a manual update as shown in the following example.

1. Go to the source appliance, and navigate to the share.

   ```
   source:shares default action-000> ls
   Properties:
                              id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                          target = target_a
                         enabled = true
                      continuous = false
                   include_snaps = true
                   max_bandwidth = unlimited
                      bytes_sent = 0
                  estimated_size = 0
             estimated_time_left = 00:00:00
              average_throughput = 0B/s
                         use_ssl = true
                     compression = on
                     export_path =
                           state = idle
               state_description = Idle (no update in progress)
                  export_pending = false
                         offline = true
                     next_update = Sync now
                       last_sync = <unknown>
                        last_try = Mon Nov 18 2019 04:40:40 GMT+0000 (UTC)
                     last_result = success
   ```

2. Start the update using `sendupdate`, and then view the status using the `ls` command.

   ```
   source:shares default action-000> sendupdate
   source:shares default action-000> ls
   Properties:
   ```

```
                           id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                       target = target1
                      enabled = true
                   continuous = false
                include_snaps = true
                max_bandwidth = unlimited
                   bytes_sent = 0
               estimated_size = 0
          estimated_time_left = 00:00:00
           average_throughput = 0B/s
                      use_ssl = true
                  compression = on
                  export_path = 
                        state = sending
            state_description = Ready (awaiting available resources to send update)
               export_pending = false
                      offline = true
                  next_update = Sync now
                    last_sync = <unknown>
                     last_try = Mon Nov 18 2019 04:40:40 GMT+0000 (UTC)
                  last_result = success

source:shares default action-000> ls
Properties:
                           id = 96366bf2-0b3c-4eec-e85b-e36e1b5bc18c
                       target = target1
                      enabled = true
                   continuous = false
                include_snaps = true
                max_bandwidth = unlimited
                   bytes_sent = 0
               estimated_size = 0
          estimated_time_left = 00:00:00
           average_throughput = 0B/s
                      use_ssl = true
                  compression = on
                  export_path = 
                        state = idle
            state_description = Idle (no update in progress)
               export_pending = false
                      offline = false
                  next_update = Sync now
                    last_sync = Mon Nov 18 2019 04:40:40 GMT+0000 (UTC)
                     last_try = Mon Nov 18 2019 04:40:40 GMT+0000 (UTC)
                  last_result = success
```

Reversing an Offline Replication (CLI)

Follow this procedure to move an offline replication package to a new local project, configured to replicate back to a source appliance.

Before You Begin

Import the offline replication stream from an NFS server to the replication target, as described in Importing a Replication Stream from the NFS Server (CLI).

1. From the replication target, navigate to the replicated package, and locate the project:

```
target:> shares replication packages
target: shares replication packages> select package-000
target:shares replication package-000> ls
Properties:
```

```
                            id = 1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
                   source_name = bigfish78
                    source_asn = d1fce51d-b8a9-6cf8-d71e-fcd4fe42cd0e
                     source_ip = 10.000.000.000:216
                   target_pool = poolA
                    replica_of = proj1
                       enabled = true
                         state = idle
             state_description = Idle (no update in progress)
                       offline = false
                   import_path =
                data_timestamp = Thu Feb 16 2017 19:10:59 GMT+0000 (UTC)
                     last_sync = Fri Feb 17 2017 03:10:11 GMT+0000 (UTC)
                      last_try = Fri Feb 17 2017 03:10:11 GMT+0000 (UTC)
                   last_result = success


Projects:
                         proj1
```

2. Enter `pkgreverse`.

```
target:shares replication package-000> pkgreverse
```

3. Optional: Set a new project name, and enable the action using the following commands:

```
target:shares replication package-000 pkgreverse> set new_project_name=new-kmm3
              new_project_name = new-kmm3
target:shares replication package-000 pkgreverse> set
enable_action_upon_reversal=true
   enable_action_upon_reversal = true
```

> **Note:**
>
> The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the appliance at the production site, the reverse operation will fail.

4. Enter `show` to confirm the properties, and then enter `commit`:

```
target:shares replication package-000 pkgreverse> show
Properties:
              new_project_name = new-kmm3
   enable_action_upon_reversal = true

host-prod:shares replication package-000 pkgreverse> commit
This action will move the contents of this package to a new local project
configured to replicate back to the source. Any data or metadata changes made
on the source since the last successful update will be lost when the new
project is replicated back to the source. If replication actions on the source
are not disabled, future updates to this package will fail.
```

5. Navigate to `shares replication actions`.

```
target:shares replication packages> cd /
target:> shares replication actions
```

6. Select the newly created action using the package ID from the previous steps. Use the package ID as `origin_pkg_id` to select the action.

```
target:shares replication actions> select origin_pkg_id=
1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
target:shares replication action-000> ls
Properties:
                            id = 6a10ce61-cc87-4850-89dd-8673f7734d03
                 origin_pkg_id = 1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
                        target = new_target
                   target_pool = p
                   source_pool = p
                replication_of = dataset1
                       enabled = true
                    continuous = false
                 include_snaps = false
   retain_user_snaps_on_target = false
                         dedup = false
  include_clone_origin_as_data = false
                 max_bandwidth = unlimited
                    bytes_sent = 0
                estimated_size = 0
           estimated_time_left = 00:00:00
            average_throughput = 0B/s
                       use_ssl = false
                   compression = on
                   export_path = 
                         state = idle
             state_description = Idle (no update in progress)
                export_pending = false
                       offline = false
                   next_update = Sync now
        replica_data_timestamp = Thu Nov 21 2019 16:17:25 GMT+0000 (UTC)
                     last_sync = <unknown>
                      last_try = <unknown>
                   last_result = <unknown>
                   replica_lag = 461:55:40
      recovery_point_objective = 
     replica_lag_warning_alert = 
       replica_lag_error_alert = 
replica_lag_over_warning_limit = false
  replica_lag_over_error_limit = false
```

7. To export the first replication update after reversal to an NFS server, enter `export_path` and the pathname of the NFS server. Enter `commit` and then enter `sendupdate`:

```
target:shares replication action-000> set export_path=nfs://nfs_server/export/
init_repl
                   export_path = nfs://nfs_server/export/init_repl (uncommitted)
target:shares replication action-000> commit
target:shares replication action-000> sendupdate
target:shares replication action-000> ls
Properties:
                            id = 6a10ce61-cc87-4850-89dd-8673f7734d03
                 origin_pkg_id = 1c0457eb-45bd-4f91-8e08-bc0dbacd40b7
                        target = new_target
                   target_pool = p
                   source_pool = p
                replication_of = dataset1
                       enabled = true
                    continuous = false
                 include_snaps = false
   retain_user_snaps_on_target = false
                         dedup = false
  include_clone_origin_as_data = false
```

```
                 max_bandwidth = unlimited
                    bytes_sent = 0
                estimated_size = 0
           estimated_time_left = 00:00:00
            average_throughput = 0B/s
                       use_ssl = false
                   compression = on
                   export_path = nfs://nfs_server/export/init_repl
                         state = idle
             state_description = Idle (no update in progress)
                export_pending = false
                       offline = false
                   next_update = Sync now
        replica_data_timestamp = Thu Nov 21 2019 16:17:25 GMT+0000 (UTC)
                     last_sync = <unknown>
                      last_try = <unknown>
                   last_result = <unknown>
                   replica_lag = 461:55:40
      recovery_point_objective =
     replica_lag_warning_alert =
       replica_lag_error_alert =
replica_lag_over_warning_limit = false
  replica_lag_over_error_limit = false
```

8. After the replication update is complete, navigate to the newly reversed package on the new target. The state description should be `idle`, as shown in the example:

```
new_target:shares replication packages> ls
Packages:

ID          STATE DATA_TIMESTAMP      SOURCE     DATASET
package-000 idle  unknown             target     <unknown>
```

9. Select the package, and enter `ls` to list its properties.

```
new_target:shares replication packages> select package-000
new_target:shares replication package-000> ls

Properties:
                            id = 6a10ce61-cc87-4850-89dd-8673f7734d03
                   source_name = target
                    source_asn = ddbd5d4e-daff-4f52-9417-cd6e893c694a
                     source_ip = 00.000.00.000:216
                   source_pool = poolA
                   target_pool = poolA
                    replica_of = <unknown>
                       enabled = true
                         state = idle
             state_description = Idle (no update in progress)
                       offline = true
                   import_path =
                data_timestamp = unknown
                     last_sync = unknown
                      last_try = unknown
                   last_result = unknown
```

10. Import the update from the NFS server.

```
new_target:shares replication package-000> set import_path=nfs://nfs_server/
export/init_repl
                   import_path = nfs://nfs_server/export/init_repl (uncommitted)
```

11. Enter `commit`, and then list the package properties to confirm the update has completed.

The property `last_result` displays `success`.

```
new_target:shares replication package-000> commit
new_target:shares replication package-000> ls
Properties:
                            id = 6a10ce61-cc87-4850-89dd-8673f7734d03
                   source_name = target
                    source_asn = ddbd5d4e-daff-4f52-9417-cd6e893c694a
                     source_ip = 00.000.00.000:216
                   source_pool = poolA
                   target_pool = poolA
                    replica_of = <unknown>
                       enabled = true
                         state = idle
             state_description = Idle (no update in progress)
                       offline = false
                   import_path =
                data_timestamp = unknown
                     last_sync = Fri Nov 22 2019 22:11:32 GMT+0000 (UTC)
                      last_try = Fri Nov 22 2019 22:11:32 GMT+0000 (UTC)
                   last_result = success
```

Related Topics

- Importing a Replication Stream from the NFS Server (CLI)
- Configuring Offline Replication (CLI)

Disabling Replication Compression (BUI)

You can disable compression when you create or edit a replication action. By default, all replication streams are compressed before being sent over the network. For more information, see Compressed Replication.

1. Source appliance: From the **Shares** menu, select **Projects**, and double-click on the project you want to edit.
2. Click the **Replication** tab.
3. Click the edit icon ✎ .
4. Click **Disable compression**, and click **Apply**.

Related Topics

- Compressed Replication
- Remote Replication Workflow

Disabling Replication Compression (CLI)

You can disable compression when you create or edit a replication action. By default, all replication streams are compressed before being sent over the network. For more information, see Compressed Replication.

1. To disable compression, navigate to the project or share and set the `compression` property to `off`, as shown in the following example:

   ```
   hostname:shares proj1 action-000> set compression=off
   ```

2. Enter `commit`, and then `show` to confirm the compression property is set to `off`.

```
hostname:shares proj1 action-000> commit
hostname:shares proj1 action-000> show
Properties:
                     id = 67f0d3d6-10af-6f30-9d4c-a60d19eb1200
             target = goby-10g
            enabled = true
         continuous = false
      include_snaps = false
      max_bandwidth = unlimited
         bytes_sent = 0
     estimated_size = 0
estimated_time_left = 00:14:35
 average_throughput = 0B/s
            use_ssl = false
        compression = off
        export_path =
              state = idle
  state_description = idle (no update in progress)
     export_pending = false
            offline = false
        next_update = Sync now
          last_sync = <unknown>
           last_try =
        last_result =
```

Related Topics

- Compressed Replication
- Remote Replication Workflow

Testing the Connection (BUI)

Use the following procedure to test the replication connection.

1. From the **Configuration** menu, select **Services**, then **Remote Replication**, then **Targets**.
2. Hover the cursor over the name of the target with the connection that you want to check, and click the edit icon ✎ .
3. In the **Edit Replication Target** dialog box, click the **Test Connection** button.

   A dialog box shows whether the certificate is trusted.
4. Verify the certificate.

   If the certificate is not trusted by the source (for example, if the source has destroyed the certificate), the dialog box presents the certificate for you to review, and you are prompted to accept or reject the certificate.

   To verify the certificate, compare the fingerprint of the displayed certificate with the fingerprint of the certificate on the target.

   a. Target appliance: From the **Configuration** menu, select **Settings**, then **Certificates**, and click the **System** tab.

   b. Click the information icon ℹ on the system default certificate.

      - If the system default certificate is automatically generated, click the information icon ℹ on the certificate that is based on the ASN UUID.
      - If the system uses a certificate that is signed by a certificate authority, click the information icon ℹ on that certificate.

   c. Compare the certificates.

Compare the certificate that is presented in the dialog box on the source appliance with the certificate that is shown in the information dialog box on the target appliance.

- If the fingerprints match, accept the certificate.
  If the fingerprints match, the presented certificate is the target's certificate and can be trusted. Accept the certificate. The certificate is added to the trust list of the source.
- If the fingerprints do not match, reject the certificate as not trusted.
  The certificate is not added to the trust list of the source.

5. Click the **Cancel** button in the **Edit Replication Target** dialog box.

Related Topics

- Creating a Replication Target (BUI)
- Editing a Replication Target (BUI)
- Replication Targets

Testing the Connection (CLI)

Use the following procedure to test the replication connection.

1. Go to `shares replication targets`.
2. Select the target with the connection that you want to check.
3. Test the connection.

   Use the `test` command to test the connection. If the certificate is trusted by the source, the message shown in the following example is presented.

   ```
   hostname:shares replication target-000> test
   Certificate is trusted
   ```

4. Verify the certificate.

   If the certificate is not trusted by the source (for example, if the source has destroyed the certificate), the certificate is presented for you to review, and you are prompted to accept or reject the certificate.

   To verify the certificate, compare the fingerprint of the displayed certificate with the fingerprint of the certificate on the target.

   a. On the target appliance, go to `configuration settings certificates system`.

   b. Select the system default certificate.

      - If the system default certificate is automatically generated, select the certificate that is based on the ASN UUID.
      - If the system uses a certificate that is signed by a certificate authority, select that certificate.

   c. Enter the `get` command.

   d. Compare the certificates.

      Compare the certificate that is presented on the source appliance with the certificate that is shown by the `get` command.

      - If the fingerprints match, accept the certificate.

If the fingerprints match, the presented certificate is the target’s certificate and can be trusted. Accept the certificate. The certificate is added to the trust list of the source.

- If the fingerprints do not match, reject the certificate as not trusted. The certificate is not added to the trust list of the source.

5. Enter the `done` command, and confirm.

Related Topics

- Creating a Replication Target (CLI)
- Editing a Replication Target (CLI)
- Replication Targets

Editing a Replication Target (BUI)

Use the following procedure to edit a replication target.

1. From the **Configuration** menu, select **Services**, then **Remote Replication**, then **Targets**.
2. Hover your cursor over the name of the target you want to edit, and click the edit icon ✎ .
3. Modify the values of the target properties.

   The **Name** and **Hostname** properties can be modified, and the **Host Match** option can be enabled or disabled. For descriptions of these properties, see Creating a Replication Target (BUI).

   An updated hostname or IP address must still resolve to the same appliance. To point to a different appliance than previously configured, create a new target to authenticate against the new appliance.
4. Click **Apply** to save any changes.

   The certificate trust check is performed. After a target is created, its certificate can become untrusted. For example, the source's administrator could remove the certificate from the list of trusted certificates, or the target's administrator could replace the certificate.

   If the certificate is not trusted by the source, the certificate is presented for you to review, and you are prompted to accept or reject the certificate as described in Testing the Connection (BUI).

Additional Tasks

- View the actions configured with the target. To modify the properties of an action, see Editing a Replication Action (BUI).
- Destroy the target, if no actions are using it.

> **Note:**
>
> A target should not be destroyed while actions are using it. Such actions will be permanently broken. The system makes a best effort to enforce this best practice but cannot guarantee that no actions exist in exported storage pools that are using a given target.

Related Topics

- Creating a Replication Target (BUI)

- Testing the Connection (BUI)
- Replication Targets
- Remote Replication Workflow
- Remote Replication Concepts
- Backing Up, Replicating, and Restoring Encrypted Projects and Shares

Editing a Replication Target (CLI)

Use the following procedure to edit a replication target.

1. Go to `shares replication targets`.
2. Select the target you want to edit.
3. Modify the values of the target properties.

   The `label` and `hostname` properties can be modified, and the value of the `host_match` option can be enabled or disabled. For descriptions of these properties, see Creating a Replication Target (CLI).

   An updated `hostname` value must still resolve to the same appliance. To point to a different appliance than previously configured, create a new target to authenticate against the new appliance.
4. Commit any changes.

   The certificate trust check is performed. After a target is created, its certificate can become untrusted. For example, the source's administrator could remove the certificate from the list of trusted certificates, or the target's administrator could replace the certificate.

   If the certificate is not trusted by the source, the certificate is presented for you to review, and you are prompted to accept or reject the certificate as described in Testing the Connection (CLI).

Additional Tasks

- View the actions configured with the target. To modify the properties of an action, see Editing a Replication Action (CLI).
- Destroy the target, if no actions are using it.

> **Note:**
>
> A target should not be destroyed while actions are using it. Such actions will be permanently broken. The system makes a best effort to enforce this best practice but cannot guarantee that no actions exist in exported storage pools that are using a given target.

Related Topics

- Creating a Replication Target (CLI)
- Testing the Connection (CLI)
- Replication Targets
- Remote Replication Workflow
- Remote Replication Concepts

- Backing Up, Replicating, and Restoring Encrypted Projects and Shares

Editing a Replication Action (BUI)

Use the following procedure to edit a replication action.

1. Navigate to the project or share, and click the **Replication** tab.
2. Select the project or share you want to edit.
3. Click the edit icon ✎ .
4. From the **Edit Replication Action** screen, modify the properties, and click **Apply**.

   For a description of replication actions, see Replication Action Properties.

Related Topics

- Editing a Replication Target (BUI)
- Replication Actions and Packages
- Remote Replication Workflow

Editing a Replication Action (CLI)

Use the following procedure to edit a replication action.

1. Navigate to `shares replication actions`, and enter `ls` to list available actions.

```
hostname:> shares replication actions
hostname:shares replication actions> ls
Actions:

ID         STATE REPLICA_DATA_TSTAMP TARGET      DATASET
action-007 idle  2019-02-17 23:01:19 targetA     berries
action-008 idle  2019-02-17 23:01:40 targetA     cherries
action-003 idle  2019-02-15 23:48:15 targetA     ocean
action-002 disbl <unknown>           targetA     oceanR
action-004 idle  <unknown>           targetA     berries
```

2. Select the action you want to edit.

```
hostname:shares replication actions> select action-007
hostname:shares replication action-007>
```

3. Modify the properties as necessary, using the `set` command.

   For a list of replication action commands for the CLI, see Replication Action Properties.

Related Topics

- Creating a Replication Action (CLI)
- Remote Replication Workflow

Monitoring Remote Replication

Use the following tasks to monitor replication progress, delays, events, and alerts. To investigate replication performance in detail, use replication analytics statistics.

- Monitoring Replication Progress - BUI, CLI
- Monitoring Replication Delays and RPO - BUI, CLI

- Replication Audit Actions
- Setting Replication Start and Finish Alerts - BUI, CLI
- Using Replication Analytics

Monitoring Replication Progress (BUI)

Use the following procedure to monitor replication progress.

1. From the **Shares** menu, select **Projects**, and select the replicated project, or select a project and then select the replicated share.
2. Click the **Replication** tab.

   The replication initial stages are displayed below the progress bar.

   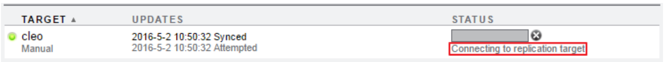

TARGET ▲	UPDATES	STATUS
cleo Manual	2016-5-2 10:50:32 Synced 2016-5-2 10:50:32 Attempted	Connecting to replication target

   The different stages are:

   - Connecting to replication target
   - Receiving checkpoint from replication target
   - Estimating size of update
   - Building deduplication tables
3. After the replication action is sending data, you can view the percentage of bytes sent, estimated size, average throughput, and estimated remaining time.

   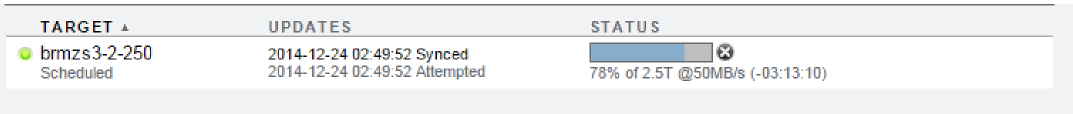

TARGET ▲	UPDATES	STATUS
brmzs3-2-250 Scheduled	2014-12-24 02:49:52 Synced 2014-12-24 02:49:52 Attempted	78% of 2.5T @50MB/s (-03:13:10)

Related Topics

- Replication Audit Actions
- Using Replication Analytics
- Deduplicated Replication

Monitoring Replication Progress (CLI)

Use the following procedure to monitor replication progress.

1. Navigate to the project or share, and enter the `replication` node.

   ```
   hostname:shares> select TestProj
   hostname:shares TestProj> replication
   hostname:shares TestProj replication>
   ```
2. Select the replication action, then enter `get`:

   ```
   hostname:shares TestProj replication> select action-000
   hostname:shares TestProj action-000> get
   Properties:
                                 id = aed46331-160b-48ec-8727-dcd563adbd78
   ```

```
                     target_id = 4fd3483e-b1f5-4bdc-9be3-b3a4becd0c42
                        target = target1
                       enabled = true
                    continuous = false
                 include_snaps = true
   retain_user_snaps_on_target = false
                         dedup = true
  include_clone_origin_as_data = false
                 max_bandwidth = unlimited
                    bytes_sent = 0
                estimated_size = 0
           estimated_time_left = 00:00:00
            average_throughput = 0B/s
                       use_ssl = true
                   compression = on
                   export_path =
                         state = sending
             state_description = Connecting to replication target
                export_pending = false
                       offline = false
                   next_update = Sync now
        replica_data_timestamp = Thu Apr 25 2019 22:18:11 GMT+0000 (UTC)
                     last_sync = <unknown>
                      last_try = <unknown>
                   last_result = <unknown>
                   replica_lag = 00:00:09
      recovery_point_objective =
     replica_lag_warning_alert =
       replica_lag_error_alert =
replica_lag_over_warning_limit = false
  replica_lag_over_error_limit = false
```

3. Review property `state_description` for information on replication progress.

   The different states are:

   - `Connecting to replication target`
   - `Receiving checkpoint from target`
   - `Estimating size of update`
   - `Building deduplication tables`
   - `Sending update`

   State `Building deduplication tables` is displayed only if the project or share has deduplication enabled.

4. If `state_description` is `Sending update`, determine the replication progress by reviewing the following properties:

   - `bytes_sent`
   - `estimated_size`
   - `estimated_time_left`
   - `average_throughput`

Related Topics

- Replication Audit Actions
- Using Replication Analytics
- Deduplicated Replication

Monitoring Replication Delays and RPO (BUI)

With asynchronous replication, a time delay occurs when writing data from the source to its replica on the target. A warning and error alert can be set to notify the administrator when a replication delay is approaching or exceeds the replication point objective (RPO). These alerts prompt the administrator to check for networking problems, application problems, and to examine the health of Oracle ZFS Storage Appliance using analytics.

A replication delay alert can be set when creating or editing a replication action. Use the following procedure to set replication delay alerts.

1. From the **Shares** menu, select **Projects**.
2. Select a project or share, and click the **Replication** tab.
3. Next to **Actions**, click the add icon ⊕ .
4. Select the properties for this action. See Replication Action Properties.
5. Select **Recovery point objective** and enter a value. Then specify **days**, **hours**, **minutes**, or **seconds**.

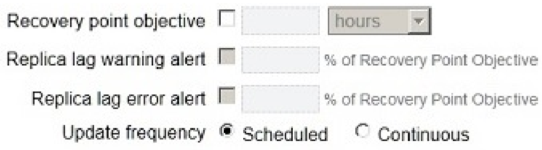

Recovery point objective □ hours ▾

Replica lag warning alert □ % of Recovery Point Objective

Replica lag error alert □ % of Recovery Point Objective

Update frequency ● Scheduled ○ Continuous

6. Select **Replica lag warning alert** and **Replica lag error alert**, and specify a percentage of the RPO for each property.

   Setting these properties will generate warnings at different times. For example, enter 50 to generate a minor alert when the replication delay exceeds 50 percent of the RPO. Enter 180 to generate a major alert when the replication delay exceeds 180 percent of the RPO.

   When the replica lag falls below the set values, a minor alert reports the replica lag is within the warning or error limit.
7. Set an alert action as described in Adding an Alert Action (BUI).

Related Topics

- Creating a Replication Action (BUI)
- Replication Action Properties

Monitoring Replication Delays and RPO (CLI)

With asynchronous replication, a time delay occurs when writing data from the source to its replica on the target. A warning and error alert can be set to notify the administrator when a replication delay is approaching or exceeds the replication point objective (RPO). These alerts prompt the administrator to check for networking problems, application problems, and to examine the health of Oracle ZS Storage Appliance using analytics.

A replication delay alert can be set when creating or editing a replication action. Use the following procedure to set replication delay alerts.

1. Navigate to the project or share, and enter `action`:

```
host_source:shares New_Project replication> action
```

2. Display the properties by entering `get`.

```
host_source:shares New_Project action (uncommitted)> get
                 origin_pkg_id =
                        target = replication-target
                          pool = demo_pool
                       enabled = true
                    continuous = false
                 include_snaps = true
   retain_user_snaps_on_target = false
                         dedup = false
  include_clone_origin_as_data = false
                 max_bandwidth = unlimited
                    bytes_sent = 0
                estimated_size = 0
           estimated_time_left = 00:00:00
            average_throughput = 0B/s
                       use_ssl = true
                   compression = on
                   export_path =
                         state = idle
             state_description = Idle (no update in progress)
                export_pending = false
                       offline = false
                   next_update = Sync now
                   replica_lag = P1H30M
        replica_data_timestamp = Wed Feb 20 2019 12:12:05 GMT+0000 (UTC)
                     last_sync = Wed Feb 20 2019 22:32:59 GMT+0000 (UTC)
                      last_try = Wed Feb 20 2019 22:32:59 GMT+0000 (UTC)
                   last_result = success
      recovery_point_objective =
     replica_lag_warning_alert =
       replica_lag_error_alert =
```

3. Set the RPO and replica lag properties for this action as shown in the following example:

```
host_source:shares New_Project action (uncommitted)> set
recovery_point_objective=50min
                         recovery_point_objective = 50 minutes (uncommitted)
host_source:shares New_Project action (uncommitted)> set
replica_lag_warning_alert=50
                             replica_lag_warning = 50% (uncommitted)
host_source:shares New_Project action (uncommitted)> set
replica_lag_error_alert=180
                     replica_lag_error = 180% (uncommitted)
```

For a description of all properties, see Replication Action Properties.

4. Commit the changes for this action.

```
host_source:shares New_Project action (uncommitted)> commit
```

5. To view the current properties for the action, enter `ls`. The RPO and replica lag portion of the output is shown in the following example:

```
host_source:shares New_Project action (uncommitted)> ls
.
.
                   replica_lag = P1H30M
        replica_data_timestamp = Wed Feb 20 2019 12:12:05 GMT+0000 (UTC)
                     last_sync = Wed Feb 20 2019 22:32:59 GMT+0000 (UTC)
                      last_try = Wed Feb 20 2019 22:32:59 GMT+0000 (UTC)
```

```
                last_result = success
    recovery_point_objective = 50 minutes
   replica_lag_warning_alert = 50%
     replica_lag_error_alert = 180%
```

6. Set an alert action as described in Adding an Alert Action (CLI).

Related Topics

- Creating a Replication Action (CLI)
- Replication Action Properties

Replication Audit Actions

The following replication configuration actions are tracked and written to the audit log. To view audit log entries in the BUI, go to **Maintenance: Logs: Audit**.

- Creating, modifying, or destroying replication actions
- Adding or removing shares from a replication group
- Creating, modifying, cloning, reversing, severing or destroying replication packages on the target
- Creating, modifying, or destroying replication targets

Setting Replication Start and Finish Alerts (BUI)

The volume of start and finish alerts for scheduled updates can obscure other important alerts. Use this procedure to disable or enable replication start and finish alerts.

Finish alerts provide statistics about each replication update. In the CLI, the `stats` node of a replication action provides statistics for the most recently completed update or for accumulated totals for this replication action.

1. From the **Configuration** menu, select **Services**, then **Remote Replication**, then **Properties**.
2. Disable or enable posting start/finish alerts.
   - To disable posting start/finish alerts, uncheck the **Enable Start/Finish Alerts** button.

     Make sure the button does not display a check mark.
   - To enable posting start/finish alerts, check the **Enable Start/Finish Alerts** button.

     Make sure the button displays a check mark.

     To view alerts, from the **Maintenance** menu, select **Logs**, then **Alerts**.
3. Click the **APPLY** button in the top-right corner of the **Properties** page.

Related Topics

- Replication Alerts
- Start and Finish Alerts
- Table "Replication Action stats Node Properties (CLI Read-Only)" in Replication Action Properties
- To specify a custom alert action, see Adding an Alert Action (BUI), and select one of the following options from the **Category** menu:
  – **Remote replication**

   - **Remote replication: source only**
   - **Remote replication: target only**

Setting Replication Start and Finish Alerts (CLI)

The volume of start and finish alerts for scheduled updates can obscure other important alerts. Use this procedure to disable or enable replication start and finish alerts.

Finish alerts provide statistics about each replication update. In the CLI, the `stats` node of a replication action provides statistics for the most recently completed update or for accumulated totals for this replication action.

1. Go to `configuration services replication`, and enter the `get` command.

   ```
   hostname:> configuration services replication
   hostname:configuration services replication> get
                              <status> = online
        enable_start_finish_alerts = true
   ```

2. Disable or enable posting start/finish alerts.

   - To disable start/finish alerts, set the value of the `enable_start_finish_alerts` property to `false`.

     ```
     hostname:configuration services replication> set
     enable_start_finish_alerts=false
          enable_start_finish_alerts = false (uncommitted)
     hostname:configuration services replication> commit
     ```

   - To enable start/finish alerts, set the value of the `enable_start_finish_alerts` property to `true`.

     To view alerts, enter `maintenance logs`, and then enter `select alert` and `show`.

3. Re-enter the `get` command to confirm that the value is correct.

Related Topics

- Replication Alerts
- Start and Finish Alerts
- Table "Replication Action stats Node Properties (CLI Read-Only)" in Replication Action Properties
- To specify a custom alert action, see Adding an Alert Action (CLI), and set the `category` property to one of the following values:
   - `replication`
   - `replication_source`
   - `replication_target`

Using Replication Analytics

The following analytics statistics are available for monitoring replication progress:

- Data Movement: Replication Bytes in *Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x*
- Data Movement: Replication Operations in *Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x*

- Data Movement: Replication Latencies (advanced analytics) in *Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x*
- Data Movement: Replication Send/Receive Bytes (advanced analytics) in *Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x*

Each statistic can be broken down by direction, type of operation, peer, pool name, project, or dataset, or can be captured as a raw statistic.

BUI - From the **Analytics** menu, select **Datasets**, and look for **Data Movement: Replication statistics**, or go to **Analytics** and add **Data Movement: Replication statistics** to a worksheet.

CLI - Go to `analytics datasets` or `analytics worksheets` and look for or add `repl.` statistics.

For more information about analytics and statistics, see Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x.

Managing Replication Packages

When projects and shares are replicated to a target, the data is stored in a replication package. Replication packages can be used for disk-to-disk backup, failover, testing purposes, or mounted from a client for read-only access.

Use any of the following methods to access data within a replication package:

- **Export a replication package** - When you export a selected replication package, the exported shares contain the data from the most recently completed replication update. After exporting the shares, you can access any specific snapshot data for any of the shares by navigating to the `.zfs/snapshot` directory in the share's root directory. The `.zfs` directory is normally invisible, but can be accessed by specifying a subdirectory or file name explicitly, as described in Accessing a Hidden Filesystem Snapshot Directory. For information about changing the snapshot visibility property, see Making a Filesystem Snapshot Directory Visible - BUI, CLI.
- **Clone a replication package** - Creating a clone of a replication package uses the most recent data received from the source appliance. Cloning converts a replication package into a new project allowing read-write access to all data in the project. The cloned project must have a unique name, mountpoint, and SMB resource name that does not conflict with any existing project. For more information, see Cloning a Replication Package - BUI, CLI.
- **Clone shares from a replication package** - Individual share snapshots, created prior to the most recent replication, can be cloned from a replication package. Cloning an individual snapshot from a replication package provides read-write access to data as it existed on the source appliance at the time a snapshot was created. For more information, see Cloning a Snapshot - BUI, CLI.
- **Sever a replication package** - The sever operation converts a replication package to a new project, allowing read-write access to data within the project. The replication connection between the source appliance and target appliance is severed after this operation. Note that the new project must have a unique name, mountpoint, and SMB resource name that does not conflict with any existing projects. For more information, see Severing a Replication Package - BUI, CLI.
- **Reverse a replication package** - Reverse replication converts a replication package into a new project, allowing read-write access to data within the project. The replication connection is preserved and a new replication action is created that allows replication back to the original source appliance. For more information, see Reverse the Direction of Replication. The multi-target reversal feature sends incremental updates from the new

source to all originally configured targets. For information on this feature, including configuration and conflict resolution, see Multi-target Reversal.

Additional tasks related to replication packages include:

- Managing User-Generated Snapshots
- Canceling a Replication Update - BUI, CLI
- Cloning a Replication Package - BUI, CLI
- Severing a Replication Package - BUI, CLI
- Editing a Replication Package - BUI, CLI
- Disabling a Replication Package - BUI, CLI

Managing User-Generated Snapshots

Disk-to-disk backup can also be achieved by setting the property `Retain user-generated snapshots on target`. User-generated snapshots, created on a source appliance, are replicated to a target which serves as an incremental-forever backup repository.

Setting this property allows you to manage user-generated snapshots independently on a source appliance and replication target. Normally, when you destroy user-generated snapshots on the source appliance, the snapshots are immediately destroyed at the replication target after a replication update. To keep user-generated snapshots on the target, set this property when creating or editing a replication action.

When user-generated snapshots are no longer needed, manually destroy them on the replication target. To destroy snapshots, see Destroying a Snapshot - BUI, CLI.

Related Topics

- Replication Action Properties
- Creating a Replication Action - BUI, CLI
- Editing a Replication Action - BUI, CLI
- Snapshot Space Management

Canceling a Replication Update (BUI)

Replication packages are displayed as projects under the **Replica** filter.

1. Target appliance: From the **Shares** menu, select **Projects**, then **Replica**.
2. Click the **Replication** tab.

   If an update is in progress, you will see a barber-pole progress bar with a cancel icon ⊗ next to it.

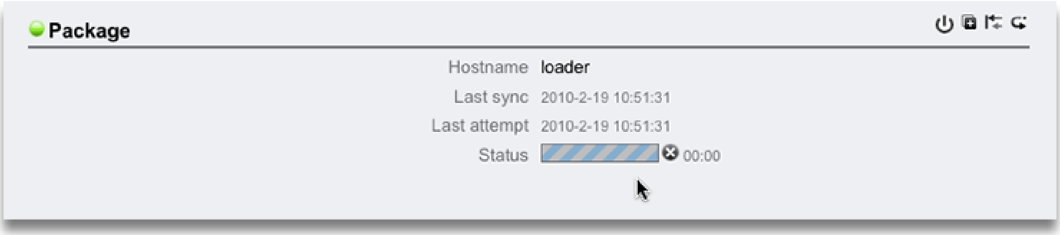

Package

Hostname loader

Last sync 2010-2-19 10:51:31

Last attempt 2010-2-19 10:51:31

Status 00:00

3. Click the cancel icon ⊗ .

   It might take several seconds for the cancellation to complete.

   After canceling an update, the next scheduled update sends the remainder of the interrupted data stream, followed by an incremental update that is based on the interrupted data stream.

   > **Note:**
   >
   > A manual update cannot be initiated from the target appliance. You must log into the source appliance to initiate a manual update.

Related Topics

- Replication Packages
- Resumable Replication
- Manually Sending a Replication Update (BUI)

Canceling a Replication Update (CLI)

You can cancel in-progress replication updates from the replication target.

1. From the replication target, navigate to `shares replication packages` and enter `ls` to list the packages.

   ```
   hostname:> shares replication packages
   hostname:shares replication packages> ls

   ID          STATE DATA TIMESTAMP      SOURCE       DATASET
   package-002 idle  2019-10-02 19:26:37 hostsource   berries
   package-001 idle  2019-10-02 19:26:10 hostsource   berries
   package-004 idle  2019-10-02 20:53:51 hostsource   berries/blackberry
   package-003 recv  2019-10-02 20:59:52 hostsource   cherries/maraschino
   ```

   Entries are sorted by `source`, `dataset`, and `data timestamp` respectively. The most recent replica snapshot is indicated by `data timestamp`.

2. Select a package.

   ```
   hostname:shares replication packages> select package-001
   ```

3. Enter `cancelupdate`.

   ```
   hostname:shares replication package-001> cancelupdate
   ```

   It might take several seconds for the cancellation to complete.

   After canceling an update, the next scheduled update sends the remainder of the interrupted data stream, followed by an incremental update that is based on the interrupted data stream.

   > **Note:**
   >
   > A manual update cannot be initiated from the replication target. You must log into the source appliance to initiate a manual update.

Related Topics

- Replication Packages
- Manually Sending a Replication Update (CLI)

Cloning a Replication Package (BUI)

A clone of a replication package is based on the most recently received replication snapshot.

When creating a cloned project, avoid naming conflicts by following these guidelines:

- The cloned project must have a unique name that does not conflict with any of the existing projects within the same pool.
- The mountpoint and SMB resource name for any of the shares of the cloned project must not conflict with any existing mountpoint or SMB resource names.
- For shares that inherit project properties, resolve conflicts by overriding the project-level mountpoint and/or SMB resource name.
- For shares that do not inherit properties from the project, set a suffix that will be appended to the mountpoint and/or SMB resource names to resolve conflicts, or override the mountpoint and/or SMB resource name individually for each share.

Use the following procedure to clone a replication package.

1. Target appliance: Navigate to the replication package you want to clone.
2. Click the **Replication** tab.
3. Click the clone icon ⊞ .
4. In the **Clone Replication Project** dialog box, complete the following fields:
   a. **New project** - Enter a unique name for the new project (clone).

      A name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters `_ - . :`

   b. Optional: **Mountpoint** - Enter a unique project-level mountpoint for the clone.

      This setting applies to shares that inherit the mountpoint from the project. Entering a unique mountpoint helps avoid conflicts.

   c. Optional: **Disable SMB Sharing** - Check to disable SMB.

      This setting applies to shares that inherit SMB sharing from the project. Shares that do not inherit SMB sharing from the project are not affected. Disabling SMB sharing on the project level will not disable SMB sharing, for shares that do not inherit SMB sharing from project.

   d. Optional: **SMB Resource Name Prefix** - Enter an SMB resource name.

      This setting applies to shares that inherit the SMB resource name from the project. Entering a unique resource name helps avoid conflicts.

      When SMB is enabled, you can share the clone over SMB. The **SMB Resource Name Prefix** used to share inherited shares of the new cloned project will be constructed using the prefix you add, plus the name of the corresponding share.

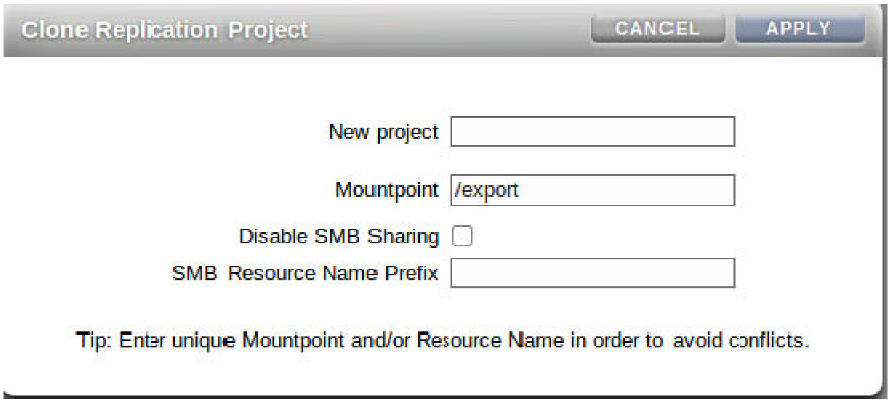

Clone Replication Project

New project

Mountpoint /export

Disable SMB Sharing □

SMB Resource Name Prefix

Tip: Enter unique Mountpoint and/or Resource Name in order to avoid conflicts.

5. Click **APPLY**.

   If there are no mountpoint or resource name conflicts, the clone operation is initiated.

   If the project name is already in use, an alert appears and a new project name must be entered.

6. Optional: If conflicts are detected, use the additional dialog boxes to resolve them.

   a. Resolve conflicts that inherit a mountpoint and/or SMB resource names from the project.

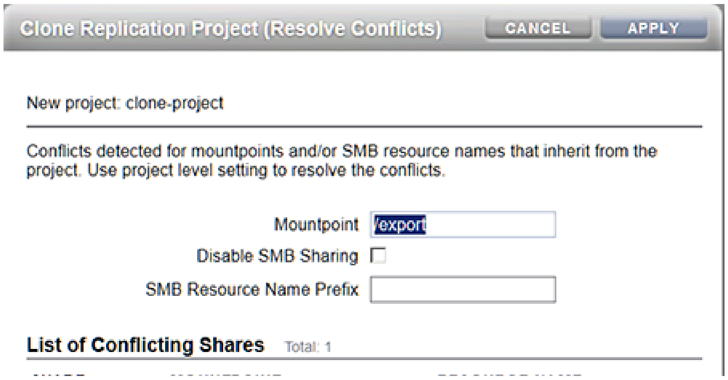

Clone Replication Project (Resolve Conflicts)

New project: clone-project

Conflicts detected for mountpoints and/or SMB resource names that inherit from the project. Use project level setting to resolve the conflicts.

Mountpoint /export

Disable SMB Sharing □

SMB Resource Name Prefix

List of Conflicting Shares Total: 1

      i. Enter a unique mountpoint.

      ii. Disable SMB Sharing or enter a unique SMB resource name prefix.

      iii. Click **APPLY**.

   b. Resolve conflicts for shares that do not inherit a mountpoint and/or SMB resource names from the project.

      This dialog box appears only after you have resolved all conflicts for inheriting shares.

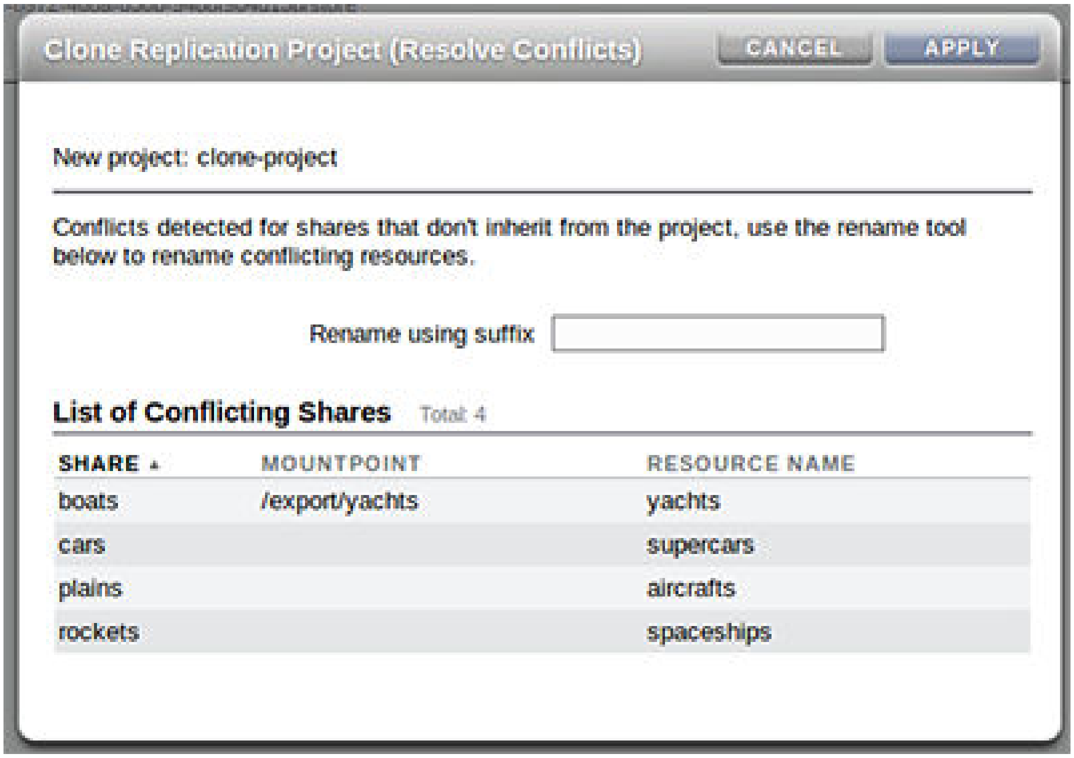

New project: clone-project

Conflicts detected for shares that don't inherit from the project, use the rename tool below to rename conflicting resources.

Rename using suffix

List of Conflicting Shares Total: 4

SHARE ▲	MOUNTPOINT	RESOURCE NAME
boats	/export/yachts	yachts
cars		supercars
plains		aircrafts
rockets		spaceships

i. Enter a unique suffix to append to the mountpoint and/or SMB resource names of these shares.

ii. Click **APPLY**.

7. If conflicts still exist, repeat step 6 to resolve the appropriate conflicts.

Related Topics

- Cloning a Replication Package or Share
- Replication Packages

Cloning a Replication Package (CLI)

A clone of a replication package is based on the most recently received replication snapshot.

When creating a cloned project, avoid naming conflicts by following these guidelines:

- The cloned project must have a unique name that does not conflict with any of the existing projects within the same pool.
- The mountpoint and SMB resource name for any of the shares of the cloned project must not conflict with any existing mountpoint or SMB resource names.
- For shares that inherit project properties, resolve conflicts by overriding the project-level mountpoint and/or SMB resource name.
- For shares that do not inherit properties from the project, set a suffix that will be appended to the mountpoint and/or SMB resource names to resolve conflicts, or override the mountpoint and/or SMB resource name individually for each share.

Use the following procedure to clone a replication package.

1. Replication target: Navigate to `shares replication packages` and enter `list` to list the packages.

```
target:> shares replication packages
target:shares replication packages> list

ID          STATE DATA TIMESTAMP      SOURCE       DATASET
package-002 idle  2019-10-02 19:26:37 hostsource   berries
package-001 idle  2019-10-02 19:26:10 hostsource   berries
package-004 idle  2019-10-02 20:53:51 hostsource   berries/blackberry
package-003 recv  2019-10-02 20:59:52 hostsource   cherries/maraschino
```

Entries are sorted by `source`, `dataset`, and `data timestamp` respectively. The most recent replica snapshot is indicated by `data timestamp`.

2. Select the replication package you want to clone.

```
target:shares replication packages> select package-001
```

3. Enter `clone` to create a new clone project.

```
target:shares replication package-001> clone
target:shares replication package-001 clone>
```

4. Set `target_project` to the project name.

   The project name must be unique, or the clone operation will fail.

   A project name must consist of 1 to 64 characters, but not include spaces or begin with a period. Allowable characters are: alphanumeric and special characters `_ - . :`

```
target:shares replication package-001 clone> set target_project=clone
        target_project = clone
```

5. Enter `conflicts` to check for conflicts.

   If conflicts exist, a message similar to the following appears.

```
target:shares replication package-001 clone> conflicts

Cloning cannot proceed because the following shares have mountpoints
or SMB resource names that are invalid or conflict with those of
other shares (either on the system or also being failed over).
Please specify valid mountpoints or SMB resource names for
these shares:

SHARE          MOUNTPOINT                      SHARESMB
share1         /export/share1                  share1
clothes        /export/clothes  (inherited)    clothes  (inherited)
electronics    /export/electronics             electronics
furniture      /export/furniture  (inherited)  furniture  (inherited)
groceries      /export/groceries  (inherited)  groceries  (inherited)
health         /export/health  (inherited)     health  (inherited)
toys           /export/toys                    toys

target:shares replication packages package-001 clone>
```

> **Note:**
>
> The `conflicts` command can be used at any point in this procedure to check for mountpoint or naming conflicts.

6. Optional: Set project-level properties to resolve conflicts for shares that inherit properties from a project.

   Use the `get` command to view the properties of the clone.

```
target:shares replication package-001 clone> get
               target_project = clone2
                rename_suffix =
          original_mountpoint = /export
                   mountpoint = /export/clone
   original_smb_resource_name = off
            smb_resource_name = off
```

The property `mountpoint` shows the current mountpoint. The property `smb_resource_name` shows the current resource name.

a. Enter a unique project-level mountpoint for the clone.

   This setting applies to shares that inherit a mountpoint from a project. Use `set mountpoint` to specify a unique mountpoint for the clone.

   ```
   target:shares replication package-001 clone> set mountpoint=/export/clone
                               mountpoint = /export/clone
   ```

b. Enter a unique project-level SMB resource name.

   This setting applies to shares that inherit the SMB resource name from a project. Set `smb_resource_name` to a unique SMB resource name.

   ```
   target:shares replication package-001 clone> set smb_resource_name=clone
                        smb_resource_name = clone
   ```

c. Set `rename_suffix` to resolve remaining share conflicts.

   This property creates a suffix that is appended to a mountpoint and/or SMB resource names, if a conflict occurs.

   ```
   target:shares replication package-001 clone> set rename_suffix=-clone
                            rename_suffix = -clone
   ```

   > **Note:**
   >
   > This operation overrides inheritance. For example, if a share originally inherits its mountpoint from the project, but the mountpoint is renamed with a suffix during the clone operation, the share in the new cloned project no longer inherits its mountpoint, but instead uses the unique renamed mountpoint.

7. Optional: To set properties for an individual share:

   a. Select a share.

      ```
      target:shares replication package-001 clone> select share1
      ```

   b. Override its mountpoint and/or SMB resource name.

      The following example overrides the mountpoint for the share.

      ```
      target:shares replication package-001 clone share1> set mountpoint=/export/
      appliances-clone
      target:shares replication package-001 clone share1> set sharesmb=appliances-
      clone
      ```

8. Enter `confirm commit` to initiate the clone operation.

   ```
   target:shares replication package-001 clone> confirm commit
   ```

   If name conflicts are detected, a message similar to the following appears:

```
Cloning cannot proceed because the following shares have
mountpoints or SMB resource names that are invalid or conflict
with those of other shares (either on the system or also being
failed over). Please specify valid mountpoints or SMB resource
names for these shares:

SHARE          MOUNTPOINT                       SHARESMB
share1         /export/share1                   share1
clothes        /export/clothes  (inherited)     clothes  (inherited)
electronics    /export/electronics              electronics
furniture      /export/furniture  (inherited)   furniture  (inherited)

target:shares replication package-001 clone>
```

9. Optional: Resolve any remaining name conflicts and confirm the cloning.

   Repeat steps 6 and 7, as appropriate, until no conflicts remain, and then enter `commit`.

Related Topics

- Cloning a Replication Package or Share
- Remote Replication Workflow

Cloning a Snapshot in a Replication Package (BUI)

> **Note:**
>
> Cloning is a licensed feature. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the appropriate Licensing Information User Manual for your product.

> **Note:**
>
> Snapshots within a replication package can be temporary. It is possible for a replication snapshot to be destroyed by replication updates, to accommodate new snapshots from the source. Therefore, this procedure recommends disabling replication updates for the package before cloning the replication snapshot.

Use the following procedure to clone a replication snapshot within a replication package.

1. Go to the package that contains the share you want to clone.
   a. From the **Shares** menu, select **Projects**.
   b. Click **Replica** above the list of projects.
   c. Hover over the package, and click the edit icon ✎ .
2. Suspend replication updates for the replication package.

   This will disable replication updates for a package entirely, which will cancel any ongoing updates and cause new updates from the source appliance to fail.
   a. In the replication package details page, select **Replication**.
   b. Click the power icon ⏻ to disable replication updates.

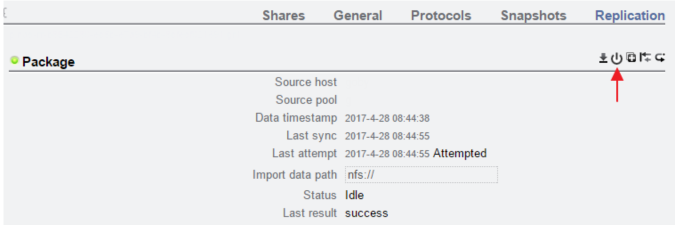

Package

Source host
Source pool
Data timestamp 2017-4-28 08:44:38
Last sync 2017-4-28 08:44:55
Last attempt 2017-4-28 08:44:55 Attempted
Import data path nfs://
Status Idle
Last result success

The status icon on the left indicates the status of the **Package**.

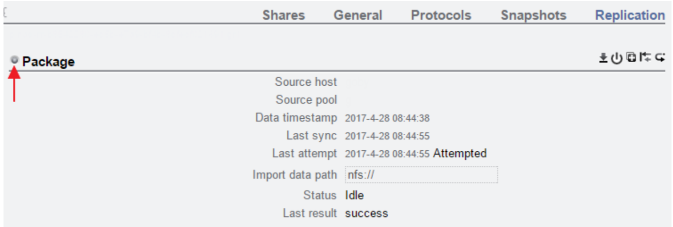

Package

Source host
Source pool
Data timestamp 2017-4-28 08:44:38
Last sync 2017-4-28 08:44:55
Last attempt 2017-4-28 08:44:55 Attempted
Import data path nfs://
Status Idle
Last result success

The **Package** remains disabled until you re-enable it using the same power icon ⏻ .

3. Navigate to the share you want to clone.

   a. In the replication package details page, select the **Shares** tab.

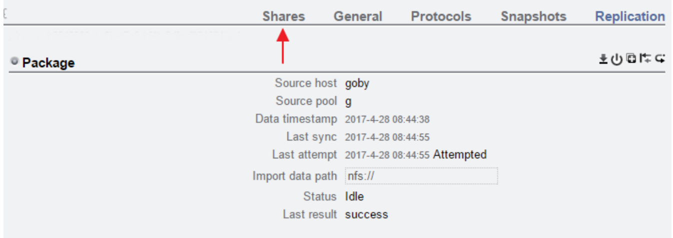

Package

Source host goby
Source pool g
Data timestamp 2017-4-28 08:44:38
Last sync 2017-4-28 08:44:55
Last attempt 2017-4-28 08:44:55 Attempted
Import data path nfs://
Status Idle
Last result success

   b. Hover over a filesystem or LUN and click the edit icon ✎ .

   c. Click the **Snapshots** tab.

4. Hover over the snapshot you want to clone, and click the clone icon ⊞ .

5. In the **Create Clone** dialog box, set the following fields.

   a. From the **Project** drop-down menu, select the destination project.

      The clone is created within the current project, by default, but you can specify a different project.

   b. Type a name for the clone.

   c. Optional: Click the lock icon 🔒 next to **Mountpoint**, and set a mountpoint for the clone.

      If you leave this field locked, the mountpoint for the clone remains as `/export/<sharename>`.

   d. Optional: Click the lock icon 🔒 next to **Resource name**, and enter one of the following values:

      - **off** - To disable SMB.
      - **on** - To enable SMB, so you can share the clone over SMB. The name of the clone in SMB matches the name of the clone in the appliance.
      - ***<resource_name>*** **-** SMB is enabled, so you can share the clone over SMB. The name of the clone in SMB is the name you specify here instead of the name of the clone in the appliance.

If you leave the **Resource** field locked, the **Resource name** property is inherited from the snapshot you are cloning.

   e. Optional: Check the **Inherit key** check box, or uncheck the check box and select the keystore and the name of the encryption key that you want the clone to inherit.

      If the box is checked, the keystore and keyname of the clone will be that of the destination project.

      If the box is unchecked, the keystore and keyname of the clone will be that of the parent share. Alternatively, select a different keystore and keyname from the drop-down menu.

   f. Optional: Check the **Retain Other Local Settings** check box to cause any inherited properties to be preserved as local settings in the new clone.

      This field determines whether inherited properties will come from the parent dataset or the destination project. By default, the box is unchecked, meaning that all inherited properties will come from the destination project for the new clone. If you check the box, all currently inherited properties will be preserved as local settings in the new clone.

6. Click **APPLY** to confirm the settings and create the clone.

   The clone appears in the list of shares for the destination project you set. You can work with a clone just like any other share.

> **Note:**
>
> After a replication package snapshot is cloned, it can no longer be destroyed by replication updates from the source.

7. Re-enable replication updates for the replication package.

   a. Navigate to the parent project of the share you just cloned.

   b. Click the **Replication** tab.

   c. Click the power icon ⏻ .

      Ensure that the status icon on the left is green, indicating that replication updates have been enabled.

Related Topics

Cloning a Replication Package or Share

Cloning a Snapshot in a Replication Package (CLI)

> **Note:**
>
> Cloning is a licensed feature. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the appropriate Licensing Information User Manual for your product.

> **Note:**
>
> Snapshots within a replication package can be temporary. It is possible for a replication snapshot to be destroyed by replication updates, to accommodate new snapshots from the source. Therefore, this procedure recommends disabling replication updates for the package before cloning the replication snapshot.

Use the following procedure to clone a replication snapshot within a replication package.

1. Go to `shares replication packages` and enter `list` to display the available replication packages.

   ```
   hostname:> shares replication packages
   hostname:shares replication packages> list
   Packages:

   ID          STATE DATA_TIMESTAMP      SOURCE     DATASET
   package-005 idle  2019-04-28 22:28:08 sor1       data1
   package-004 idle  2019-04-28 15:44:38 sor1       data1
   package-003 disbl 2019-04-27 23:46:20 sor1       data1
   package-002 idle  2019-04-27 23:14:10 sor1       data1
   package-001 idle  2019-04-17 17:27:05 sor2       data2
   ```

2. Select the package that contains the share you want to clone.

   ```
   hostname:shares replication packages> select package-005
   ```

3. Suspend replication updates for the replication package.

   This action disables replication updates for a package entirely, which cancels any ongoing updates and causes new updates from the source appliance to fail.

   a. Enter `set enabled=false`.

      ```
      hostname:shares replication package-005> set enabled=false
                           enabled = false (uncommitted)
      ```

   b. Enter `commit`.

      ```
      hostname:shares replication package-005> commit
      ```

4. Select the project that contains the share you want to clone.

   a. Enter `show` to display the project name.

      ```
      hostname:shares replication package-005> show
      Properties:
                                  id = 7e184188-2738-432b-f304-123412341234de
                                  ...
                                  ...
                                  ...

      Projects:
                                 proj1
      ```

   b. Select the project.

      ```
      hostname:shares replication package-005> select proj1
      ```

5. Select the share you want to clone.

   a. Enter `show` to display the available shares.

```
hostname:shares replication package-005 proj1> show
Properties:
                    aclinherit = restricted
                             ...
                             ...
                             ...

Shares:

Filesystems:

NAME             SIZE   ENCRYPTED     MOUNTPOINT
share1           36K    off           /export/share1
share2           36K    off           /export/share2
```

b. Select the share.

```
hostname:shares replication package-005 proj1> select share1
```

6. Select the snapshot you want to use to clone the share.

   a. Enter `snapshots`.

```
hostname:shares replication package-005 proj1/share1> snapshots
```

   b. Enter `list` to display the available snapshots.

```
hostname:shares replication package-005 proj1/share1 snapshots> list
Snapshots:
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-cb
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-ec
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f2
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f3
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f4
.rr-e1401958-9f7b-47bf-8245-fa116972d26f-f5

Children:
                        automatic => Configure automatic snapshots
```

   c. Select the snapshot you want to clone.

```
hostname:shares replication package-005 proj1/share1 snapshots> select .rr-e1401958-9f7b-47bf-8245-fa116972d26f-cb
```

7. Clone the snapshot.

   a. Use the `clone` command, followed by the name of the project in which you want to create the clone, and the name for the clone.

```
hostname:shares replication package-005 proj1/share1@.rr-e1401958-9f7b-47bf-8245-fa116972d26f-cb> clone proj_name clone1
```

   b. Use the `get` command to view properties.

```
hostname:shares proj_name/clone1 (uncommitted clone)> get
Properties:
                    aclinherit = restricted (inherited)
                       aclmode = discard (inherited)
                         atime = true (inherited)
                      checksum = fletcher4 (inherited)
                   compression = off (inherited)
                        copies = 1 (inherited)
                       logbias = latency (inherited)
                    mountpoint = /export/clone1 (inherited)
                         quota = 0 (default)
                      readonly = false (inherited)
```

```
                                ...
                                ...
                                ...
```

   c. Use the `set` command to adjust properties.

      ```
      hostname:shares proj_name/clone1 (uncommitted clone)> set mountpoint=/export/
      clone_mountpoint_name
                          mountpoint = /export/clone_mountpoint_name (uncommitted)
      ```

   d. Use the `commit` command to commit the changes and create the clone.

      ```
      hostname:shares proj_name/clone1 (uncommitted clone)> commit
      hostname:shares replication package-005 proj1/share1@.rr-e1401958-9f7b-47bf-8245-
      fa116972d26f-cb>
      ```

8. Re-enable replication updates for the replication package.

   a. Enter `cd /` to return to the top level.

      ```
      hostname:shares replication package-005 proj1/share1@.rr-e1401958-9f7b-47bf-8245-
      fa116972d26f-cb> cd /
      ```

   b. Enter `shares replication`, then use `select` and the package name.

      ```
      hostname:> shares replication
      hostname:shares replication> select package-005
      ```

   c. Use `set enabled=true` to re-enable replication updates to the package. Then enter `commit` to save the changes.

      ```
      hostname:shares replication package-005> set enabled=true
                             enabled = true (uncommitted)
      hostname:shares replication package-005> commit
      ```

9. Check your specified project destination to see the clone.

   a. Enter `cd /` to return to the top level.

      ```
      hostname:shares replication package-005> cd /
      ```

   b. Use `shares select` with the project you used for the clone destination.

      ```
      hostname:> shares select proj_name
      ```

   c. Enter `show` to list the shares, and look for the cloned share.

      ```
      hostname:shares proj_name> show
      Properties:
                          aclinherit = restricted
                             aclmode = discard
                                    ...
                                    ...
                                    ...

      Shares:

      Filesystems:

      NAME                     SIZE   ENCRYPTED     MOUNTPOINT
      clone1                   1K     off           /export/clone_mountpoint_name
      ```

> **Note:**
>
> After a replication package snapshot is cloned, it can no longer be destroyed by replication updates from the source.

Related Topics

Cloning a Replication Package or Share

Severing a Replication Package (BUI)

Use the following procedure to sever a replication package.

Before You Begin

Check for mountpoint or SMB share name conflicts between replicated filesystems and other filesystems on the system. To resolve mountpoint (and/or SMB resource names) conflicts, change the project or share mountpoints (or SMB resource names) in the replication package before severing a replication package. For more information, see Severing Replication.

1. Navigate to the replication package.
2. Click the **Replication** tab.
3. Click the sever icon .
4. Enter a name for the new local project.

> **Note:**
>
> If a replication update is performed during or after a sever operation, the update fails with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target.

Related Topics

- Severing Replication
- Remote Replication Concepts

Severing a Replication Package (CLI)

Use the following procedure to sever a replication package.

Before You Begin

Check for mountpoint or SMB share name conflicts between replicated filesystems and other filesystems on the system. To resolve mountpoint (and/or SMB resource names) conflicts, change the project or share mountpoints (or SMB resource names) in the replication package before severing a replication package. For more information, see Severing Replication.

1. From the replication target, navigate to the replication package.

   ```
   host-target:shares default replication source-001 package-001>
   ```

2. Enter `sever` and a name for the new local project.

   ```
   host-target:shares default replication source-001 package-001> sever new_project
   ```

> **Note:**
>
> If a replication update is performed during or after a sever operation, the update fails with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target.

Related Topics

- Severing Replication
- Remote Replication Concepts

Editing a Replication Package (BUI)

Use the following procedure to edit a replication package.

1. Target appliance: From the **Shares** menu, select **Projects**, then **Replica**.

   The name, size, and creation date of each replication package is displayed.

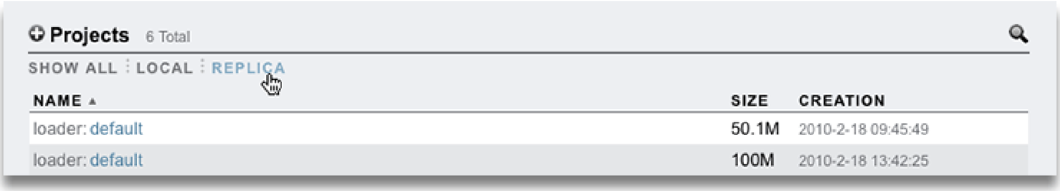

NAME ▴	SIZE	CREATION
loader: default	50.1M	2010-2-18 09:45:49
loader: default	100M	2010-2-18 13:42:25

> **Note:**
>
> Packages are displayed in the BUI only after the first replication update has begun. They may not appear in the list until some time after the first update has completed.

2. Select a replication package for editing.

   The **Shares** view for the package's project is displayed.

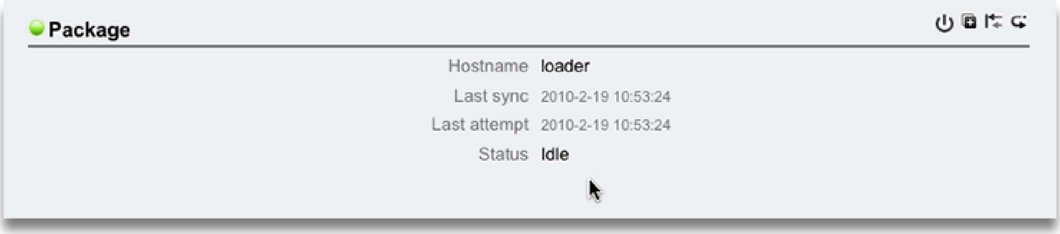

Hostname loader
Last sync 2010-2-19 10:53:24
Last attempt 2010-2-19 10:53:24
Status Idle

3. To modify package properties, click the **Replication** tab.

   For a list of properties that you can modify, see Replication Packages.

Related Topics

- Replication Actions and Packages

- Remote Replication Concepts

Editing a Replication Package (CLI)

Replication packages are organized in the CLI in a flat view under `shares replication packages`, which displays all replication packages in the system.

Use the following procedure to edit a replication package.

1. From the replication target, go to `shares replication packages`, and enter `list` to list all replication packages in the system.

```
hostname:> shares replication packages
hostname:shares replication packages> list

ID          STATE  DATA TIMESTAMP      SOURCE       DATASET
package-002 idle  2019-10-04 19:26:37 hostsource   berries
package-001 idle  2019-10-04 19:26:10 hostsource   berries
package-004 idle  2019-10-04 20:53:51 hostsource   berries/blackberry
package-003 recv  2019-10-04 20:59:52 hostsource   cherries/maraschino
```

   The packages are sorted by `SOURCE`, `DATASET`, and `DATA_TIMESTAMP` (in descending order).

2. Select a package.

```
hostname:shares replication packages> select package-001
hostname:shares replication packages package-001> show

Properties:
                id = d6137c89-7056-4788-a4d1-b5892fe315e0
       source_name = hostsource
        source_asn = a751dc0f-abcd-1234-6789-f5e8315eaffa
         source_ip = 00.000.00.00:000
       source_pool = poolS
       target_pool = poolT
        replica_of = berries
           enabled = true
             state = idle
 state_description = Idle (no update in progress)
           offline = false
       import_path =
    data_timestamp = Fri Oct 04 2019 19:26:10 GMT+0000 (UTC)
         last_sync = Fri Oct 04 2019 19:26:10 GMT+0000 (UTC)
          last_try = Fri Oct 04 2019 19:26:10 GMT+0000 (UTC)
       last_result = success

Projects:
           berries
```

   A replication package can be selected directly by specifying its ID, as shown in the following example:

```
hostname:shares replication packages> select d6137c89-7056-4788-a4d1-
b5892fe315e0
hostname:shares replication packages package-001>
```

3. To edit project properties and shares, select a project.

```
hostname:shares replication packages package-001> select berries
hostname:shares replication packages package-001 berries> get mountpoint
```

```
                    mountpoint = /export
hostname:shares replication packages package-001 berries> get sharenfs
                      sharenfs = on
```

For a list of package properties that you can modify, see Replication Package Properties.

Related Topics

- Replication Packages
- Remote Replication Concepts

Disabling a Replication Package (BUI)

Replication updates for a package can be disabled entirely, which will cancel any ongoing updates and cause new updates from the source appliance to fail.

1. Target appliance: Navigate to the package.
2. Click the **Replication** tab.
3. Click the power icon ⏻.

   The status icon on the left indicates the status of the package (enabled, disabled, or failed). The package remains disabled until you re-enable it using the using the same power icon ⏻.

Related Topics

- Replication Packages
- Remote Replication

Disabling a Replication Package (CLI)

Replication updates for a package can be disabled entirely, which will cancel any ongoing updates and cause new updates from the source appliance to fail.

1. Target appliance: Navigate to the package.
2. Modify the `enabled` property.
3. Commit your changes.

   The package remains disabled until you set the `enabled` property to `on`.

Related Topics

- Replication Packages
- Remote Replication Concepts

Disaster Recovery with Remote Replication

A two-system disaster recovery site consists of a source appliance at a production site and a replication target located at a recovery site in a geographically different location. In the event of a catastrophic production site failure, the administrator redirects client operations to the recovery site by reversing replication on the replication target, thus ensuring continuous operation. After the production site is restored to normal operation, the administrator updates the production site by reversing replication at the recovery site. To restore the original source-target relationship, replication is then reversed again.

Note:

Replication reversal is not supported to a filesystem with mandatory file retention. For information about the file retention feature, see File Retention Management.

To set up remote replication for disaster recovery, use these tasks:

- Setting Up a Target Appliance at a Recovery Site - BUI, CLI
- Switching Operations to the Recovery Site - BUI, CLI
- Updating the Production Site - BUI, CLI
- Reversing Replication Back to the Production Site - BUI, CLI

Setting Up a Replication Target at a Recovery Site (BUI)

Use the following procedure to create a replication target for disaster recovery.

1. Identify a replication target at a recovery site.

   The replication target requires a software version compatible with the source appliance. For details, see MOS DOC ID 1958039.1 (https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1).
2. Source appliance: Create a target as described in Creating a Replication Target (BUI).
3. Create a replication action, and schedule a continuous replication. See Creating a Replication Action (BUI).

   Note:

   Continuous replication minimizes data loss in the event of a disaster at the production site.

Next Steps

Switching Operations to the Recovery Site (BUI)

Switching Operations to the Recovery Site (BUI)

After a failure occurs at the production site, perform a reverse replication at the recovery site, and then redirect client operations to the recovery site.

1. Replication target: From the **Shares** menu, select **Projects**, then **Replica**, and look for the replication package from the source appliance.

   The replica is named *source_appliance*:*project*/*share*.
2. Double-click the replication package, or click its edit icon ✎ .
3. Click the **Replication** tab.
4. Click the reverse replication direction icon ⇆ .
5. Enter a name for the new local project, and enable the action.

   The project name and location of the original action are preserved.

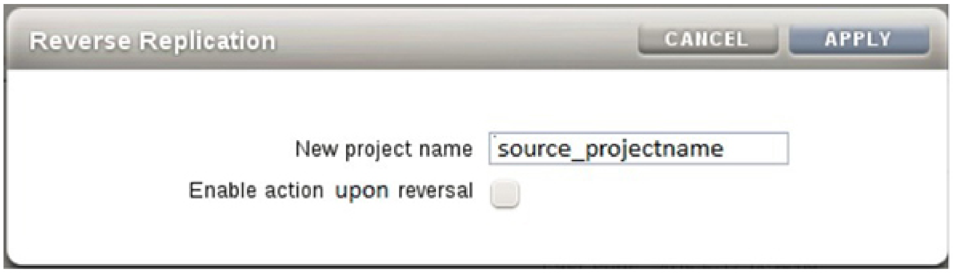

New project name: source_projectname

Enable action upon reversal □

> **Note:**
>
> The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the target at the recovery site, the reverse operation will fail.

6. Click **OK**.
7. Transfer client activity to the IP address of the appliance at the recovery site.

   Depending on the protocol used, map (SMB clients) or remount (NFS clients) the shares using the IP address or name of the appliance at the recovery site.

Next Steps

Updating the Production Site (BUI)

Updating the Production Site (BUI)

After the production site is restored and back online, replicate the changes written to the recovery site during the outage back to the production site.

1. Appliance at the recovery site: From the **Shares** menu, select **Projects**, then **Local**, and select the new local project.

   The new project is listed with status `Never synced`.
2. Click **Sync Now** to start the replication.
3. Wait for the replication to complete.

   At the top of the window, `Finished replicating to the new` *`project`* `on the` *`source`* `appliance` is displayed.

Next Steps

Reversing Replication Back to the Production Site (BUI)

Reversing Replication Back to the Production Site (BUI)

After all changes have been replicated from the recovery site to the production site, use reverse replication again to restore the original replication relationship between source and replication targets. To send incremental updates from the new source to multiple, originally configured targets, see Multi-target Reversal.

1. Production appliance: From the **Shares** menu, select **Projects**, then **Replica**, and look for the new project name.

   The project is named *target_appliance*: *new_project*/*share*.

2. Select the new project, and click its edit icon ✎ .
3. Click the **Replication** tab.
4. Click the reverse replication direction icon ⇄ .
5. In the **Reverse Replication** window, enter a name for the new local project, and then enable the action.

   > **Note:**
   >
   > The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the source at the production site, the reverse operation will fail.

6. Depending on the protocol used, remap (SMB clients) or remount (NFS clients) shares to the appliance at the recovery site.
7. Delete the original project on the source appliance.
   a. From the **Shares** menu, select **Projects**, then **Local**, and look for the original project, which should be empty.
   b. Select the empty project, and click its destroy icon 🗑 .
   c. Click **OK**.

Setting Up a Replication Target at a Recovery Site (CLI)

Use the following procedure to create a replication target for disaster recovery.

1. Identify a replication target at the recovery site.

   The replication target requires a software version compatible with the source appliance. For details, see MOS DOC ID 1958039.1 (https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1).
2. Source appliance: Create a target as described in Creating a Replication Target (CLI).
3. Create a replication action, and schedule a continuous replication. See Creating a Replication Action (CLI).

   > **Note:**
   >
   > Continuous replication minimizes data loss in the event of a disaster at the production site.

Next Steps

Switching Operations to the Recovery Site (CLI)

Switching Operations to the Recovery Site (CLI)

After a failure occurs at the production site, perform a reverse replication at the recovery site, and then redirect client operations to the recovery site.

1. Replication target: Enter `shares replication`.

   ```
   host-offsite:> shares replication
   ```

2. Enter `sources` to list the source appliances that are associated with this target.

   ```
   host-offsite:shares default replication> sources
   ```

3. Look for the package replicated by the source appliance.

   In this example, `source-001` is the source appliance number, and `host-prod` is the source appliance name. `kmm2` is the local project name in replication package number `package-001`.

   ```
   source-001 host-prod
    PROJECT                    STATE         LAST UPDATE
   package-000 <unknown>        idle          unknown
   package-001 kmm2             idle          Wed May 01 2019 20:06:27 GMT+0000(UTC)
   ```

4. Enter `select` and the source appliance number.

   ```
   host-offsite:shares default replication sources> select source-001
   ```

5. Enter `select` and the replication package number.

   ```
   host-offsite:shares default replication source-001> select package-001
   ```

6. Enter `pkgreverse`.

   ```
   host-prod:shares replication source-005 package-000> pkgreverse
   ```

   The `pkgreverse` command preserves the properties of the original replication action, including schedules.

7. Optional: Set a new project name, and enable the action using the following commands:

   ```
   host-prod:shares replication source-000 package-000 pkgreverse> set
   new_project_name=new-kmm2
                 new_project_name = new-kmm2
   host-prod:shares replication source-000 package-000 pkgreverse> set
   enable_action_upon_reversal=true
      enable_action_upon_reversal = true
   ```

   > **Note:**
   >
   > The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the target at the recovery site, the reverse operation will fail.

8. Enter `show` to confirm the properties, and then enter `commit`:

   ```
   host-prod:shares replication source-000 package-000 pkgreverse> show
   Properties:
                 new_project_name = new-kmm2
      enable_action_upon_reversal = true

   host-prod:shares replication source-000 package-000 pkgreverse> commit
   This action will move the contents of this package to a new local project
   configured to replicate back to the source. Any data or metadata changes made
   on the source since the last successful update will be lost when the new
   project is replicated back to the source. If replication actions on the source
   are not disabled, future updates to this package will fail.
   ```

9. Transfer client activity to the IP address of the appliance at the recovery site.

   Depending on the protocol used, map (SMB clients) or remount (NFS clients) the shares using the IP address or name of the appliance at the recovery site.

Next Steps

Updating the Production Site (CLI)

Updating the Production Site (CLI)

After the production site is restored and back online, replicate the changes written to the recovery site during the outage to the production site.

1. Appliance at the recovery site: Go to `shares` and select the new project.

   ```
   host-offsite:> shares
   host-offsite:shares pool> select new-kmm2
   ```

2. Enter `list` to find the name of the share.

   ```
   host-offsite:shares pool new-kmm2> list
   Filesystems:
   NAME              SIZE  ENCRYPTED     MOUNTPOINT
   karen2            31K    off           /export/karen2
   host:shares pool new-kmm2> replication
   host:shares pool new-kmm2 replication> show
   Actions:
   TARGET         STATUS     NEXT
   action-000
   host2          idle       Sync now
   ```

3. Select the action number, and then enter `sendupdate` to start replication to the production appliance.

   ```
   host-offsite:shares pool new-kmm2 replication> select action-000
   host-offsite:shares pool new-kmm2 action-000> sendupdate
   ```

4. Wait for the replication to complete.

   The state changes to `idle` when the replication has completed.

Next Steps

Reversing Replication Back to the Production Site (CLI)

Reversing Replication Back to the Production Site (CLI)

After all changes have been replicated from the recovery site to the production site, reverse replication again to restore the original replication relationship between source and replication targets. To send incremental updates from the new source to multiple, originally configured targets, see Multi-target Reversal.

1. Replication target: Navigate to `shares replication packages` and enter `list` to list the packages.

   ```
   loader:> shares replication packages
   loader:shares replication packages> list

   ID          STATE DATA TIMESTAMP      SOURCE       DATASET
   package-002 idle  2019-10-02 19:26:37 hostsource   berries
   package-001 idle  2019-10-02 19:26:10 hostsource   berries
   package-004 idle  2019-10-02 20:53:51 hostsource   berries/blackberry
   package-003 recv  2019-10-02 20:59:52 hostsource   cherries/maraschino
   ```

   The package with the most recent `data timestamp` for this dataset is the one with the most recent, up-to-date data for the corresponding source dataset.

2. Select the replication package you want to reverse.

   ```
   loader:shares replication packages> select package-002
   ```

3. Enter `pkgreverse`.

   ```
   host-prod:shares replication package-002> pkgreverse
   ```

4. Optional: Set a new project name, and enable the action using the following commands:

   ```
   host-prod:shares replication package-002 pkgreverse> set new_project_name=new-
   kmm3
                 new_project_name = new-kmm3
   host-prod:shares replication package-002 pkgreverse> set
   enable_action_upon_reversal=true
      enable_action_upon_reversal = true
   ```

   > **Note:**
   >
   > The project name must be unique on the appliance where the reverse operation is performed. If a project with the same name exists on the appliance at the production site, the reverse operation will fail.

5. Enter `show` to confirm the properties, and then enter `commit`:

   ```
   host-prod:shares replication package-002 pkgreverse> show
   Properties:
                 new_project_name = new-kmm3
      enable_action_upon_reversal = true

   host-prod:shares replication package-002 pkgreverse> commit
   This action will move the contents of this package to a new local project
   configured to replicate back to the source. Any data or metadata changes made
   on the source since the last successful update will be lost when the new
   project is replicated back to the source. If replication actions on the source
   are not disabled, future updates to this package will fail.
   ```

6. Depending on the protocol used, remap (SMB clients) or remount (NFS clients) shares to the appliance at the production site.

Related Topics

- Reverse the Direction of Replication
- Remote Replication Concepts

Remote Replication Concepts

Oracle ZFS Storage Appliance Replication is a licensed feature for certain models that provides snapshot-based replication of projects and shares from a source appliance to one or more replication targets. Replication performs a full update of an entire project and/or share contents, followed by incremental updates containing only the changes since the previous update.

For license details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

This topic describes key remote replication (or simply "replication") concepts and replication terminology.

Remote replication has the following important characteristics:

- **Snapshot-based asynchronous replication** - The replication subsystem takes a snapshot as part of each update operation. In the case of an initial update, the entire contents of a project and/or share are sent. In the case of an incremental update, only the changes since the last replication snapshot for the same action are sent. Because replication takes snapshots and then sends them, data is already committed to stable storage before replication even begins sending it. Continuous replication effectively sends continuous streams of filesystem changes, but it is still asynchronous with respect to NAS and SAN clients.
- **Block-level** - Each update operation traverses the filesystem at the block level and sends the appropriate filesystem data and metadata to the target.
- **Includes all metadata** - The underlying replication stream serializes both user data and metadata, including most properties configured on the Shares screen. These properties can be modified on the target after the first replication update completes, though not all take effect until the replication connection is severed. For example, this allows sharing over NFS to a different set of hosts than on the source. For more information, see Replication Packages.
- **Secure** - Secure Sockets Layer (SSL) is always used to secure the replication control protocol used among Oracle ZFS Storage Appliance products and optionally replication data transfers. The appliance can only replicate to or from another appliance after an initial manual authentication process. For more information, see Replication Targets.
- **Encrypted projects and shares** - When enabled, transparent data encryption protects individual shares (filesystems and LUNs) and projects. For more information, see Data Encryption.
- **Protocol independent** - The appliance supports both file and block-based storage volumes. The replication mechanism is protocol independent.
- **Compressed replication** - Support for compressed replication streams increases replication performance and improves throughput utilization between multiple sites that have limited bandwidth connections. For more information, see Compressed Replication.

Replication has the following known limitations:

- Actions cannot move between pools.
- Network throughput is limited to a maximum of 200 MB/s per project level replication. With compressed replication, the effective data rate can exceed the actual physical network data rate.

For information about remote replication concepts, see the following topics:

- Replication Terminology
- Replication Targets
- Replication Actions and Packages
- Replication Action Properties
- Replication Storage Pools
- Project vs. Share Replication
- Replication Authorizations
- Deduplicated Replication
- Replication Configuration for Clustered Appliances
- Example: Replication Configuration for Clustered Appliances

- Replication Snapshots and Data Consistency
- Replication Snapshot Management
- iSCSI Configurations and Replication
- Resumable Replication
- Replication Alerts
- Replication Failures
- Compressed Replication
- Replication Packages
- Cloning a Replication Package or Share
- Multi-target Reversal
- Cascaded Replication
- Exporting Replicated Filesystems
- Severing Replication
- Reverse the Direction of Replication
- Destroying a Replication Package
- Target Replica Backups
- Raw Crypto Replication

Replication Terminology

The following is a list of the common replication terms.

- **clone** - A replicated package can be cloned into a mutable project. A clone can be managed like any other project on the system.
- **full sync or full update** - A replication operation that sends the entire contents of a project and some of its shares. The initial synchronization of a project and/or share is a full synchronization.
- **incremental update** - A replication operation that sends only the differences in a project and its shares since the previous update (whether that one was full or incremental).
- **recovery point objective (RPO)** - The maximum tolerable amount of data loss, expressed in unit of time, in the event of a disaster or major outage. The RPO is defined as part of a disaster recovery plan and represents the last consistent set of data that is available for recovery. The lower an RPO value, the less data loss.
- **replica** - Replicated data contained in the replication package on the replication target.
- **replication action** - Describes the data to be replicated, the replication schedule, and data transfer properties such as enabling or disabling encryption of the network link. See Replication Action Properties.
- **replication group** - The set of datasets (exactly one project and some number of shares) that are replicated as a unit. See Project vs. Share Replication.
- **replication package** - Exists on the replication target and is associated with a replication action. It is a special object that contains the replica. A replication package can be exported, cloned, severed, or reversed, which allows for write access to the data within the project.

- **replication source** - An appliance that sends replication updates to one or more target appliances, periodically, continuously, or on demand. Individual appliances can act as both a source and a *target*, but are only one of these in the context of a particular replication *action*.
- **replication target** - An appliance that receives and stores data replicated from one or more *source* appliances. This term also refers to a configuration object on the appliance that enables it to replicate to another appliance.
- **reverse replication** - A replication relationship that exchanges the source and target roles. Following a disaster recovery, roles can be reversed again.

Replication Targets

A replication target can be:

- A different pool on the source appliance.
- A separate NFS server for offline replication.

When you create a replication target, first enter a name for the target. This name is used to identify the target in the BUI and CLI of the source appliance.

Establishing a Secure Connection

The replication target and the source appliance establish a connection that enables secure communications between the source and target. Provide the following information to establish the connection:

- Fully qualified domain name, or IPv4 or IPv6 address of the target appliance - The recommended value to use is the target's domain name.
- Root password of the replication target - Authorizes the administrator to set up the connection on the replication target.

The source and target exchange keys that are used to securely identify each other in subsequent communications. These keys are stored as part of the appliance configuration and persist across reboots and upgrades. These keys are lost if the appliance is factory reset or reinstalled.

The root password is never stored persistently. Changing the root password on either the target or source does not require any change to the replication configuration. The password is never transmitted in the clear. The initial identity exchange - and all other replication operations - is protected with SSL.

If you need to ensure that the replication traffic goes over a particular network interface, set up a static route for the target that specifies that interface as shown in Setting Up Network Interfaces and Static Routing - BUI, CLI.

Verifying the Target Certificate

When you create a replication target, certificate verification is performed. Certificate verification consists of the following steps:

1. Certificate hostname is checked.
2. Certificate trust is checked.

If either the hostname check or the trust check fails, the target is not created.

Hostname Check

The value of the hostname property can be a fully qualified domain name, or an IPv4 or IPv6 address. The recommended value to use is the target's fully qualified domain name.

The hostname check verifies that the hostname that you specified for the target matches a host specified in the certificate. If you specify an IP address or an unqualified domain name for the hostname, and the certificate only has fully qualified domain names, the hostname check fails and the target is not created.

If the target is using an ASN-based certificate, specify the target's fully qualified domain name for the value of the hostname property.

The hostname check is performed by default. You can bypass the hostname check by disabling the host match option.

For stronger security, set the value of the hostname property to the target's fully qualified domain name, and make sure the host match option is enabled.

Certificate Trust Check

The certificate trust check verifies that one of the following certificates has been added to the source's trusted certificate list and is enabled for peer use:

- The target appliance's certificate.
- The certificate for the certificate authority that issued the target appliance's certificate.

The certificate is verified when you create or edit a target, and any time the source and target try to connect. You can also check the connection yourself at any time.

- **Create target** - When you add a target, if the certificate is not trusted by the source, the certificate is presented for you to review, and you are prompted to accept or reject the certificate. If you accept the certificate, the certificate is added to the trust list of the source, and the target is created. If you reject the certificate, the certificate is not added to the trust list of the source, and the target is not created. If the certificate is already trusted, the target is created, and you are not prompted to accept the certificate.

  After the target is created, its certificate can become untrusted. For example, the source's administrator could remove the certificate from the list of trusted certificates, or the target's administrator could replace the certificate.
- **Edit target** - When you apply changes, if the certificate is not trusted by the source, the certificate is presented for you to review, and you are prompted to accept or reject the certificate. If you accept the certificate, the certificate is added to the trust list of the source. If you reject the certificate, the certificate is not added to the trust list of the source.
- **Connection test** - At any time, you can check whether the certificate is trusted. If the certificate is not trusted by the source, the certificate is presented for you to review, and you are prompted to accept or reject the certificate. See Testing the Connection - BUI, CLI.
- **Peer and replication connection** - The certificate trust check is performed for every peer and replication connection. If the certificate is not trusted, the source rejects the connection.

Related Topics

- Creating a Replication Target - BUI, CLI
- Testing the Connection - BUI, CLI
- Editing a Replication Target - BUI, CLI

Other Replication Target Considerations

If a replication source uses NIS or LDAP and directly maps these service's users or groups in the share configuration, the equivalent setup must be present on the replication target. Otherwise, replication sever and reverse operations could fail.

By default, the replication target connection is not bidirectional. For example, if an administrator configures replication from source `S` to target `T`, `T` cannot automatically use `S` as a target. However, reversing the direction of replication is supported, which automatically creates a target for `S` on `T` (if it does not already exist) so that `T` can replicate back to `S`. For more information, see Reverse the Direction of Replication.

Related Topics

- Creating a Replication Target - BUI, CLI.
- Remote Replication Workflow
- Remote Replication Concepts

Replication Actions and Packages

A replication action specifies where and how to replicate a project or share. Replication actions are created on the source appliance and specify the following parameters:

- A replication group that includes either a project or individual shares
- Name of the replication target
- Name of a storage pool on the replication target (used only during the initial setup)
- Frequency (scheduled or continuous) of the update
- Number of Auto-Snapshots (Scheduled Snapshots) that are retained on the target
- Additional options such as encryption of the data stream or disabling compression

A replication group is specified implicitly by the project or share on which the action is configured (see Replication Storage Pools). The replication target and storage pool cannot be changed after the action is created, but the other options can be modified at any time. If a replication update is in progress when an option is changed, then the new value takes effect when the next update begins (with the exception of the `max bandwidth` parameter, which takes effect immediately after the modification).

When a replication action is created on the source appliance, a package on the target in the specified storage pool is created. The package on the replication target contains an exact copy of the source project and shares on which the action is configured as of the start time of the last replication update. Actions are the primary unit of replication configuration on the appliance.

Replication Update Frequency

Replications can be executed manually or configured in the replication action to be sent continuously or at scheduled times. The three replication modes are:

- **Manual** - Replication is started manually, at any time, by the administrator. A manual replication update can be useful for testing purposes and for applications that require data to be in a certain state before replication can occur. See Manually Sending a Replication Update - BUI, CLI.

- **Scheduled** - Replication is automatically executed according to a selected schedule. The scheduled frequency can be set to replicate to the target every 5, 10, 15, 20 or 30 minutes, every 1, 2, 4, 8 or 12 hours, every day, every week, or every month. More granular update frequencies can be set by defining multiple schedules for a single replication action.

  The **Auto** selection, which is available when creating the first schedule for a replication action, is a start time generated by the appliance. When multiple replication actions are configured on an appliance, the auto-generated start time can minimize overlapping replication updates and improve load balancing.

  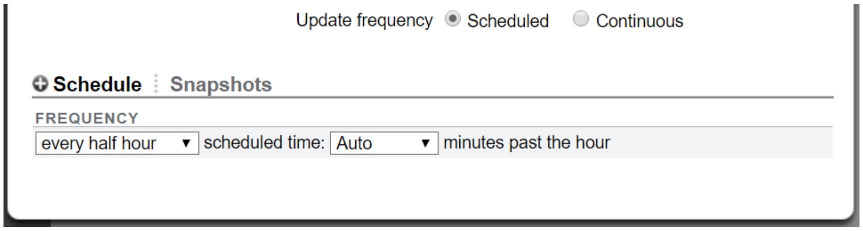

  Update frequency ● Scheduled ○ Continuous

  Schedule | Snapshots

  FREQUENCY

  every half hour ▾ scheduled time: Auto ▾ minutes past the hour

  The scheduled frequency can also be set to replicate to the target based on the automatic snapshot schedules configured in the project or share. When this option is selected, replication updates are performed when the scheduled automatic snapshots are created.

  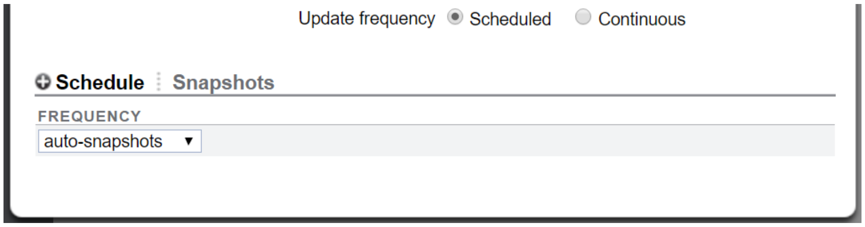

  Update frequency ● Scheduled ○ Continuous

  Schedule | Snapshots

  FREQUENCY

  auto-snapshots ▾

- **Continuous** - Replication is executed continuously. As soon as one replication update is complete, a subsequent update is started. Changes are transmitted as frequently as possible, resulting in sending a constant stream of all filesystem changes to the target system. For filesystems with a lot of churn (many files created and destroyed in short intervals), this can result in replicating much more data than is actually necessary. However, as long as replication can keep up with the data changes, this results in the minimum amount of data lost in the event of a data-loss disaster on the source system.

Replication Action and Package Relationship

A replication action and package are bound to each other. If the package is somehow corrupted or destroyed, the action cannot send replication updates, even if the target still has the data and snapshots associated with the action. Similarly, if the action is destroyed, the package will be unable to receive new replication updates (even if the source still has the same data and snapshots). A warning will occur, in both the BUI and CLI, if you attempt to perform an operation that would destroy the action-package connection. If an error or explicit administrative operation breaks the action-package connection such that an incremental update is no longer possible, you must sever or destroy the package and action, then create a new action on the source.

Related Topics

- Replication Action Properties
- Replication Packages

Replication Action Properties

The replication action properties in the BUI and CLI are different, as described in the following tables.

> **Note:**
>
> When you change a replication action property, the new setting takes effect with the next replication update unless specified otherwise.

The following table describes replication action properties that are available in both the BUI and the CLI.

Table 7-1 Replication Action Properties (BUI and CLI)

BUI Property	CLI Property	Description
Disable compression	`compression`	The replication stream is compressed by default. Disable if compression is provided by another means, such as a WAN accelerator. For more information, see Compressed Replication.
Disable raw crypto mode	`rawcrypto`	Raw crypto replication is enabled by default, and data is not decrypted on the source appliance before it is sent to the target appliance. To disable in the CLI, set `rawcrypto` to `off`. For more information, see Raw Crypto Replication.
Distant target	`distant_target`	In a multi-target reversal and cascaded replication configuration, it defines the topology transformation during the reverse on the potential source. Can be set for only project-level, non-cascading actions when `potential_target` is `TRUE`. For more information, see Cascaded Replication.
Update frequency	`continuous`	In the BUI, select **Scheduled** or **Continuous** for **Update** frequency. In the CLI, set the `continuous` property to `true` or `false`. For more information, see Replication Update Frequency.
Enable deduplication	`dedup`	When set, enables deduplication on replication streams. For more information, see Deduplicated Replication.

Table 7-1 (Cont.) Replication Action Properties (BUI and CLI)

BUI Property	CLI Property	Description
Power icon ⏻	`enabled`	When enabled (`true` in the CLI), replication updates can be sent. When disabled (`false` in the CLI), the power icon ⏻ is not highlighted, and replication updates cannot be sent.
Export data path	`export_path`	Specifies the path to an NFS share for this action, using the format: nfs://*server*/*path*. This property exports the replication stream to a file on an NFS server, which can be physically moved to the remote target site, and then imported to a replication target. For procedures, see Creating an Offline Replication - BUI, CLI.
Include clone origin as data	`include_clone_origin_as_data`	Controls the replication of each share that was cloned from a share that is external to the replication package on the target. Select this option to insert a complete copy of the clone origin snapshot's data into the replica of the clone. If you deselect this option, a clone created from an external origin snapshot will share storage with the replica of the clone origin snapshot that resides in replication target's pool. Sharing storage saves space, but if the replication target does not contain the external clone origin snapshot, the replication of the clone will fail. For details, see Cloning a Replication Package or Share.
Include Snapshots	`include_snaps`	Determines whether replication updates include non-replication snapshots. For details, see Replication Snapshot Management.
Limit bandwidth	`max_bandwidth`	Specifies a maximum speed for this replication update (in terms of amount of data transferred over the network per second). Use this property to limit bandwidth, especially during an initial replication update. If you change this property during a replication update, the new setting takes effect immediately.

Table 7-1 (Cont.) Replication Action Properties (BUI and CLI)

BUI Property	CLI Property	Description
Pool	`pool`	Storage pool on the target where this project will be replicated. This property is specified when an action is initially configured and not shown thereafter.
Potential source	`potential_source`	When multi-target reversal is configured, the corresponding target can perform a reverse of the package. Can be set for only project-level, non-cascading actions. For more information, see Multi-target Reversal.
Recovery point objective □ *unit of time*	`recovery_point_objective`	Specifies the maximum tolerable amount of data loss in the event of a disaster or major outage. The recovery point objective (RPO) can be specified as days, hours, minutes, or seconds. Set this property when creating or editing a replication action. This property is only used in conjunction with the replica lag warning and error alerts.
Replica lag error alert □ % of Recovery Point Objective	`replica_lag_error_alert`	Specifies a limit represented as a percentage of the RPO. The source appliance generates a **major alert** when the replica lag exceeds the specified limit. When the replica lag falls below this value, a minor alert indicates the replica lag is within the error limit.
Replica lag warning alert □ % of Recovery Point Objective	`replica_lag_warning_alert`	Specifies a limit represented as a percentage of the RPO. The source appliance generates a **minor alert** when the replica lag exceeds the specified limit. When the replica lag falls below this value, a minor alert indicates the replica lag is within the warning limit.
Retain user snapshots on target	`retain_user_snaps_on_target`	Retains user-generated snapshots on the replication target for the associated action, even after the snapshots are destroyed on the source appliance. The snapshots retained on the target can then be used as part of a disk-to-disk or data backup solution. For details, see Managing User-Generated Snapshots.
Target	`target`	Unique identifier for the replication target system. This property is specified when an action is initially configured and cannot be edited.

Table 7-1 (Cont.) Replication Action Properties (BUI and CLI)

BUI Property	CLI Property	Description
Target package ID	`target_pkgid`	Read-Only. ID of the replication package on the target appliance. Starting with OS8.8.0, `package id` and `action id` might not match in some cases.
Target Pool	`target_pool`	Storage pool on the target where this project will be replicated. This property is displayed when editing an existing replication action.
Enable SSL-encryption	`use_ssl`	When set, encrypts data on the wire using SSL. Using this feature may have an impact on per-action replication performance.
Update cascade delay	`update_cascade_delay`	Update delay when used with cascading replication schedule option after_update. In BUI, under **Cascading Schedule** when **After Update** is selected. For more information, see Cascaded Replication.
Schedule	`schedule`	Schedules a replication update to start at a specified minute, hour, day, week, or month. Frequency values are 5, 10, 15, 20 or 30 minutes, 1, 2, 4, 8 or 12 hours, every day, every week, or every month. If more granular update frequencies are needed, define multiple schedules for a single replication action. Auto-snapshots is an auto-generated start time selected by the appliance. Auto-snapshots is available only when one schedule is set for an action.
Snapshots	`autosnaps`	**BUI** - The **Snapshots** subtab in a replication action displays configured automatic snapshot schedules, and allows you to change the number of snapshots retained on the target. **CLI** - Automatic snapshot retention settings are configured through child node `autosnaps`. See the following table for more information. For information about automatic snapshot management, see Replication Automatic Snapshot Management.

The following table describes replication action properties that are available only in the CLI. All of these properties are read-only.

Table 7-2 Replication Action Properties (CLI Read-only)

CLI Property	Description
`average_throughput`	Describes the average replication throughput.
`bypassed_id`	ID of the bypassed cascading replication action. Enables locating a bypassed action. Available only after performing the retarget/bypass procedure as described in Cascaded Replication.
`bytes_sent`	Number of bytes sent to the replication target.
`estimated_size`	Estimated size of the data to be replicated.
`estimated_time_left`	Estimated time remaining until completion of the replication update.
`export_pending`	Indicates whether an export is pending. Value is `true` or `false`.
`id`	The replication action's unique identifier. The identifier can be used to select an action or its associated replication package using the syntax: `select id=<uniqueid>`.
`next_update`	Date and time when the next attempt will be made. This value could be a date (for a scheduled update), `Sync now`, or `continuous`.
`offline`	Indicates whether the replication update is offline. Value is `true` or `false`.
`origin_pkg_id`	The unique identifier of the replication package from which this action was created when the package was reversed. The origin package identifier is displayed after the first successful replication update completes.
`replica_lag`	Current replica lag with a format of `hh:mm:ss`.
`replica_lag_over_error_limit`	`true` when the replica lag has exceeded the error limit specified by the combination of the `recovery_point_objective` and the `replica_lag_error_alert` threshold.
`replica_lag_over_warning_limit`	`true` when the replica lag has exceeded the warning limit specified by the combination of the `recovery_point_objective` and the `replica_lag_warning_alert` threshold.
`replication_of`	Project or share name (under project) where the action is configured.
`source_pool`	Pool name of the source project/share.
`state`	Describes if replication is in progress or not. Values are `sending` or `idle`.

Table 7-2 (Cont.) Replication Action Properties (CLI Read-only)

CLI Property	Description
state_description	Specifies details on replication progress. The different states are: • Connecting to replication target • Receiving checkpoint from target • Estimating size of update • Building deduplication tables • Sending update State Building deduplication tables is displayed only if the project or share has deduplication enabled.
target_id	The unique identifier of the replication target object that describes the target of this replication action.

The following table describes nodes of a replication action. These nodes are only viewable in the CLI. Properties of these nodes are read-only.

Table 7-3 Replication Action Child Nodes (CLI only)

Replication Action Child Node	Description
autosnaps	Automatically scheduled snapshots, with sub-node automatic. To change the number of snapshots retained for a replication package, select an individual automatic node, and modify the keep property. For information about automatic snapshot management, see Replication Automatic Snapshot Management. For more information about configuring auto snapshots, see Configuring Automatic Snapshot Retention on a Target (CLI).
schedules	Schedule configuration. Allows configuration of source and cascading replication schedules. For more information, see Cascaded Replication.
stats	Statistics for the most recent replication update, and the accumulated statistics over the lifetime of this replication action. Statistics are updated after replication update completion.

The following table describes the stats node properties. The stats node properties report statistics from the most recent replication update (last) and the accumulated statistics over the lifetime of this replication action (total). To see the statistics for a replication update that is older than the most recent update, see the associated finish alert as described in Start and Finish Alerts. For more information about properties with dedup and dd_ in their names, see Deduplicated Replication.

Table 7-4 Replication Action stats Node Properties (CLI Read-Only)

CLI Property	Description
replica_data_timestamp	Creation time of the snapshot used in the most recently completed (last_sync) replication update.

Table 7-4 (Cont.) Replication Action stats Node Properties (CLI Read-Only)

CLI Property	Description
`last_result`	The result of the `last_sync` update. The value of `last_result` is either `success` or `failed`.
`last_sync`	The last time a replication update was successfully sent. This value might be unknown if the system has not sent a successful update since boot.
`last_try`	The last time a replication update was attempted. This value might be unknown if the system has not attempted to send an update since boot. If the value of `last_result` is `success`, then `last_sync` and `last_try` are the same time value.
`last_logical_bytes`	Number of bytes that the `last_sync` update data stream would have contained if the data on disk had not been compressed and without any subsequent compression or deduplication.
`last_phys_bytes`	Number of bytes in the `last_sync` internal replication update data stream prior to replication deduplication or replication data stream compression.
`last_after_dedup`	Number of bytes in the `last_sync` internal replication update data stream after any deduplication of the replication data stream.
`last_to_network`	Number of bytes that the `last_sync` update data stream compression pipeline delivered to the network. This value shows the consequence of replication data stream compression, if enabled.
`last_duration`	Elapsed time to perform the `last_sync` update.
`last_dd_table_build`	Time to build the `last_sync` update deduplication tables prior to the transmission of the replication update.
`last_dd_table_mem`	Maximum amount of memory that was consumed by the `last_sync` update deduplication tables.
`total_updates`	Number of successful replication updates for this action.
`total_logical_bytes`	Number of bytes that update data streams of all updates for this action would have contained if the data on disk had not been compressed and without any subsequent compression or deduplication.
`total_phys_bytes`	Number of bytes in internal replication update data streams of all updates for this action prior to replication deduplication or replication data stream compression.
`total_after_dedup`	Number of bytes in internal replication update data streams of all updates for this action after any deduplication of the replication data stream.
`total_to_network`	Number of bytes that update data stream compression pipelines delivered to the network for all updates for this action. This value shows the consequence of replication data stream compression, if enabled.
`total_duration`	Elapsed time to perform all updates for this action.
`dd_total_updates`	Number of successful deduplicated replication updates for this action.

Table 7-4 (Cont.) Replication Action stats Node Properties (CLI Read-Only)

CLI Property	Description
`dd_total_logical_bytes`	Number of bytes that update data streams of all deduplicated updates for this action would have contained if the data on disk had not been compressed and without any subsequent compression or deduplication.
`dd_total_phys_bytes`	Number of bytes in internal replication update data streams of all deduplicated updates for this action prior to replication deduplication or replication data stream compression.
`dd_total_after_dedup`	Number of bytes in internal replication update data streams of all deduplicated updates for this action after deduplication of the replication data stream.
`dd_total_to_network`	Number of bytes that update data stream compression pipelines delivered to the network for all deduplicated updates for this action. If replication data stream compression is enabled, `dd_total_to_network` shows how well the compression worked.
`dd_total_duration`	Elapsed time to perform all deduplicated updates for this action.
`dd_total_table_build`	Time to build the deduplication tables prior to the transmission of the replication update for all deduplicated updates for this action.
`dd_total_table_mem`	Maximum amount of memory that was consumed by the deduplication tables for all deduplicated updates for this action.

Related Topics

- Creating a Replication Action (BUI)
- Creating a Replication Action (CLI)

Replication Package Properties

The replication package properties in the BUI and CLI differ slightly, as described in the following table.

Table 7-5 Replication Package Properties (BUI and CLI)

BUI Property	CLI Property	Description
Source host	`source_name`	Name of this package's source.
Source pool	`source_pool`	Storage pool on the source where this project is replicated from. This property is specified when an action is initially configured.
⏻	`enabled`	When enabled (`true` in the CLI), replication updates can be received. When disabled (`false` in the CLI), the power icon ⏻ is not highlighted, and replication updates cannot be received.
Data timestamp	`data_timestamp`	Creation time of the snapshot used in the last successful update.
Last sync	`last_sync`	Completion time of last successful update.

Table 7-5 (Cont.) Replication Package Properties (BUI and CLI)

BUI Property	CLI Property	Description
Last attempt	`last_try`	Completion time of last update attempt.
Import data path	`import_path`	External media URI for pending import.
Status	`state`	Current state of replication updates.
Last result	`last_result`	Result of last update attempt. Value is `success` or `failed`.

The following table describes the CLI read-only replication package properties.

Table 7-6 Replication Package Properties (CLI Read-only)

CLI Property	Description
`id`	Unique identifier of the replication package. The identifier can be used to select a package, by entering `select id=` *`unique-id`*.
`source_asn`	ID of the replication action associated with this package.
`source_ip`	IP address of this package's source.
`target_pool`	Target pool for this package.
`replica_of`	Replicated dataset in this package.
`state_description`	Current state of replication updates.
`offline`	Indicates that the package is waiting for an offline update.

Replication Storage Pools

When a replication action is initially configured, the administrator is given a choice of which storage pool on the target should contain the replicated data. The storage pool containing an action cannot be changed once the action has been created. Creating the action creates the empty package on the target in the specified storage pool. After this operation the source has no knowledge of the storage configuration on the target. It does not keep track of which pool the action is being replicated to, nor is it updated with storage configuration changes on the target.

When the target is a clustered system, the chosen storage pool must be one owned by same controller which owns the IP address used by the source for replication because only those pools are always guaranteed to be accessible when the source contacts the target using that IP address. This is exactly analogous to the configuration of NAS clients, for example, NFS or SMB, where the IP address and path requested in a mount operation must follow the same constraint. When performing operations that change the ownership of storage pools and IP addresses in a cluster, administrators must consider the impact to sources replicating to the cluster. There is currently no way to move packages from one storage pool to another.

Related Topics

- Remote Replication Workflow
- Remote Replication

Project vs. Share Replication

Although it is possible to configure remote replication on both the project level and share level, project-level replication is recommended for the following reasons:

- Replication snapshots are always taken at the project level. Multiple share level replications in a single project can create a substantial amount of overhead and consume space on the pool.
- When reversing share-level replication, the share is placed in its own project. This means that replication reversals will end up splitting the share away from the other shares in the project, unless they are all replicated together.

Like other properties configurable on the **Shares** screen, each share can either inherit or override the configuration of its parent project. Inheriting the configuration means not only that the share is replicated on the same schedule to the same target with the same options as its parent project is, but also that the share will be replicated in the same stream using the same project-level snapshots as other shares inheriting the project's configuration. This may be important for applications which require consistency between data stored on multiple shares. Overriding the configuration means that the share will not be replicated with any project-level actions, though it may be replicated with its own share-level actions that will include the project. It is not possible to override part of the project's replication configuration and inherit the rest.

More precisely, the replication configuration of a project and its shares define some number of replication *groups*, each of which is replicated with a single stream using snapshots taken simultaneously. All groups contain the project itself (which essentially just includes its properties). One project-level group includes all shares inheriting the replication configuration of the parent project. Any share that overrides the project's configuration forms a new group consisting of only the project and the share itself.

For example, suppose you have:

- A project `home` and shares `bill`, `cindi`, and `dave`.
- `home` has replication configured with some number of actions.
- `home/bill` and `home/cindi` inherit the project's replication configuration.
- `home/dave` overrides the project's replication configuration, using its own configuration with some number of actions.

This configuration defines the following replication groups, each of which is replicated as a single stream per action using snapshots taken simultaneously on the project and shares:

- One project-level group, including `home`, `home/bill`, and `home/cindi`.
- One share-level group, including `home` and `home/dave`.

> **Note:**
>
> Due to current limitations, do not mix project- and share-level replications within the same project. This avoids unpredictable results when reversing the replication direction or when replicating clones. For more details, see Replication Packages and Cloning a Replication Package or Share.

Related Topics

- Remote Replication Workflow
- Remote Replication

Replication Authorizations

The replication subsystem provides two user authorizations under the "Projects and Shares" scope.

Authorization	Details
`rrsource`	Allows administrators to create, edit, and destroy replication targets and actions. Additionally, it allows an administrator to send and cancel updates for replication actions.
`rrtarget`	Allows administrators to manage replicated packages, including disabling replication at the package level, cloning a package or its members, modifying properties of received datasets, and severing or reversing replication. Other authorizations may be required for some of these operations (like setting properties or cloning individual shares). See the available authorizations in the Projects and Shares scope for details.

The `rrsource` authorization is required to configure replication targets on an appliance, even though this is configured under the **Remote Replication** service screen. For help with authorizations, see Configuring Users.

Related Topics

- Remote Replication Workflow
- Remote Replication

Deduplicated Replication

Deduplicated replication provides the ability to reduce the amount of data sent over the network by replication jobs. This feature is useful for reducing the on-the-wire data bandwidth requirements of replication, especially when using a high-latency, low-bandwidth, high-cost network.

> **Note:**
>
> This feature imposes a cost in the form of pre-processing and increased memory overhead. The effectiveness of deduplication is highly data dependent, so it is strongly recommended to verify the deduplication savings with representative datasets prior to using this feature in a production environment. Deduplicated replication is more efficient when there is more duplicate data.

Deduplicated replication is disabled by default. To enable deduplicated replication for individual replication actions, click Enable deduplication in the **Add Replication Action** dialog box in the BUI, or set the `dedup` property to `true` in the CLI.

Deduplicated Replication Statistics

In the CLI, each replication update action has a `stats` node. The `stats` node records information about the most recent replication update, as well as the accumulated statistics over the lifetime of the replication action. To view statistics for a specific update that was not the most recent update, see the finish alert for the update as described in Start and Finish Alerts.

These replication action `stats` node properties quantify:

- On-disk compression benefits
- Deduplication benefits
- Replication data stream compression benefits
- Replication update duration
- Deduplication tables construction time (before sending data)
- Deduplication tables maximum memory consumption

Table "Replication Action stats Node Properties (CLI Read-Only)" in Replication Action Properties describes the action `stats` node properties of a deduplicated replication stream. See especially properties with `dedup` and `dd_` in their names.

Measuring Deduplicated Replication Statistics

When deduplication is enabled for a replication stream, the data is transformed through several layers of deduplication and compression. Data rates are measured and recorded as the data is transformed.

To determine whether deduplication was effective for the replication action, examine the replication statistics in the `stats` node of a replication action in the CLI, or in finish alerts in the BUI or the CLI.

Single Deduplicated Replication Update Benefits Comparison

- In the BUI, use the replication finish alerts to compare the `phys_bytes` and `after_dedup` statistics to evaluate the benefit of deduplicated replication. For information about replication finish alerts, see Start and Finish Alerts.
- In the CLI, use the replication finish alerts to compare the `phys_bytes` and `after_dedup` statistics or use the replication action `stats` node to compare `last_phys_bytes` and `last_after_dedup` statistics to evaluate the benefit of deduplicated replication. For information about statistics in the `stats` node, see table "Replication Action stats Node Properties (CLI Read-Only)" in Replication Action Properties.

Averaged Deduplicated Replications Updates Benefits Comparison

To determine the average benefit of all deduplicated replication updates performed by this replication action, use the replication action `stats` node to compare statistics `dd_total_phys_bytes` and `dd_total_after_dedup`. For information about statistics in the `stats` node, see table "Replication Action stats Node Properties (CLI Read-Only)" in Replication Action Properties.

Replication Configuration for Clustered Appliances

Replication can be configured from any source appliance to any replication target regardless of whether each is part of a cluster and whether the appliance's cluster peer has replication configured in either direction.

The following rules govern the behavior of replication updates for clustered appliances:

- Replication updates for projects and shares are sent from whichever cluster peer has imported the containing storage pool.
- Storage pools that are owned by each controller can replicate to the same replication target, only if software version OS8.6.0 or later is installed on both controllers.
- Replication updates are received by whichever peer has imported the IP address configured in the replication action on the source. Administrators must ensure that the controller using this IP address will always have the storage pool containing the replica imported. This is ensured by assigning the pool and IP address resources to the same controller during cluster configuration.
- Replication updates (both to and from an appliance) that are in progress when an appliance exports the corresponding storage pool or IP address (as part of a takeover or failback) will fail. Replication updates using storage pools and IP addresses unaffected by a takeover or failback operation will be unaffected by the operation.

Related Topics

Example: Replication Configuration for Clustered Appliances

Example: Replication Configuration for Clustered Appliances

The goal of this example is to configure replication properly to ensure that projects continue to replicate after a cluster takeover, cluster failback, or after performing reverse replication on a target appliance.

- Configuration Guidelines
- Example: Configuring Replication for Clustered Appliances
- Replication Data Path Illustrated Examples

Configuration Guidelines

When configuring replication for clustered Oracle ZFS Storage Appliance systems, follow these guidelines:

- Ensure that both replication source and target appliances are in the `CLUSTERED` state. For details, see table "Cluster States" in Clustered Controller States.
- Select network interfaces and IP addresses to be used for replication traffic on the replication source and target appliances.
  - Select a singleton network interface. Unlike a private network interface, a singleton network interface will be taken over by the surviving controller following the loss of one of the controllers in the cluster. Using a singleton interface ensures successful replication following a cluster takeover or failback transition. For more information about singleton interfaces, see table "Cluster Resource Types" in Cluster Resource Management.
  - Ensure that the selected network interface on the source appliance and the pool from which the data will be replicated are both assigned to the same controller. This is always the case when the source cluster is in the `CLUSTERED` state.
  - Similarly for the target cluster, the selected network interface on the target appliance and the pool into which the data will be replicated must both be assigned to the same controller. This association is guaranteed when the replication configuration is performed while the target cluster is in `CLUSTERED` state.
  - The source and the target appliances must be able to successfully communicate using the selected network interfaces and IP addresses.

- Create static /32 host-based routing between target and source appliances to ensure that following replication reversal, the selected interface is used for outbound replication traffic when reversal has transformed the current target into a replication source.
- After the static route has been created, configure the replication target object on the source appliance using the selected IP address of the target.
- When the target appliance is in the `OWNER` state, all shared resources including network interfaces and storage pools are taken over and owned by the one surviving controller, the controller that is now in the `OWNER` state. On the controller in the `OWNER` state, it is possible to select a network interface that is assigned to one controller and use it to deliver replication traffic to a pool that is assigned to a different controller. When the controllers are returned to the `CLUSTERED` state, the network interfaces and storage pools are returned to their assigned controllers. Therefore, replication updates might not be possible because the source appliance will use the network interface on the target controller that no longer owns the pool. This configuration error cannot arise when replication configuration is performed while the target appliance is in the `CLUSTERED` state.

Example: Configuring Replication for Clustered Appliances

The example procedure uses the following source and target network interfaces and IP addresses:

The source appliance cluster consists of source controllers `S1` and `S2`. Storage pool `sp1` is assigned to `S1` and pool `sp2` is assigned to `S2`. The cluster network interfaces consist of:

- Private interface `ixgbe0` on `S1` with IP address 198.51.100.81/24
- Private interface `ixgbe0` on `S2` with IP address 198.51.100.82/24
- Singleton interface `ixgbe1` with IP address 192.0.2.101/25 assigned to `S1`
- Singleton interface `ixgbe2` with IP address 192.0.2.102/25 assigned to `S2`
- Singleton interface `ixgbe3` with IP address 192.0.2.201/25 assigned to `S1`
- Singleton interface `ixgbe4` with IP address 192.0.2.202/25 assigned to `S2`

The appliance is Initially in the `CLUSTERED` state where:

- `S1` owns `sp1`, `ixgbe1`, and `ixgbe3`
- `S2` owns `sp2`, `ixgbe2` and `ixgbe4`

The target appliance cluster consists of controllers `T1` and `T2`. Storage pool `tp1` is assigned to `T1` and pool `tp2` is assigned to `T2`. The cluster network interfaces consist of:

- Private interface `ixgbe0` on `T1` with IP address 198.51.100.83/24
- Private interface `ixgbe0` on `T2` with IP address 198.51.100.84/24
- Singleton interface `ixgbe1` with IP address 192.0.2.103/25 assigned to `T1`
- Singleton interface `ixgbe2` with IP address 192.0.2.104/25 assigned to `T2`
- Singleton interface `ixgbe3` with IP address 192.0.2.203/25 assigned to `T1`
- Singleton interface `ixgbe4` with IP address 192.0.2.204/25 assigned to `T2`

The appliance is initially in the `CLUSTERED` state where:

- `T1` owns `tp1`, `ixgbe1`, `ixgbe3`
- `T2` owns `tp2`, `ixgbe2` and `ixgbe4`

The following steps describe how to configure replication using the CLI for projects Red, Blue, and Green.

1. Select network interfaces and IP addresses.
   - Start by selecting network interfaces and IP addresses for replication of project Red.

     Because the source S is in the CLUSTERED state, it is sufficient to ensure that the selected network interfaces and IP addresses are not private. Thus, on S1 use either ixgbe1 or ixgbe3.
   - The same applies to target T, therefore, use either ixgbe1 or ixgbe3 on appliance T1. Because ixgbe1 and ixgbe3 on both S1 and T1 belong to the same subnet, select either to perform replication of project Red. For this example, select interface ixgbe1 on S1 and on T1.
2. Set up a static route on S1.

   The following example sets up the static route for replication of project Red on source controller S1:

```
S1:configuration net routing> create
S1:configuration net route (uncommitted)> set family=IPv4
                            family = IPv4 (uncommitted)
S1:configuration net route (uncommitted)> set destination=192.0.2.103
                       destination = 192.0.2.103 (uncommitted)
S1:configuration net route (uncommitted)> set mask=32
                              mask = 32 (uncommitted)
S1:configuration net route (uncommitted)> set interface=ixgbe1
                         interface = ixgbe1 (uncommitted)
S1:configuration net route (uncommitted)> set gateway=192.0.2.1
                           gateway = 192.0.2.1 (uncommitted)
S1:configuration net route (uncommitted)> commit
S1:configuration net routing> list
ROUTE      DESTINATION      GATEWAY    INTERFACE  TYPE    STATUS
...
route-003  192.0.2.103/32   192.0.2.1  ixgbe1     static  active
```

3. Set up a static route on T1.

   The following example sets the static route for replicating project Red on target controller T1:

```
T1:configuration net routing> create
T1:configuration net route (uncommitted)> set family=IPv4
                            family = IPv4 (uncommitted)
T1:configuration net route (uncommitted)> set destination=192.0.2.101
                       destination = 192.0.2.101 (uncommitted)
T1:configuration net route (uncommitted)> set mask=32
                              mask = 32 (uncommitted)
T1:configuration net route (uncommitted)> set interface=ixgbe1
                         interface = ixgbe1 (uncommitted)
T1:configuration net route (uncommitted)> set gateway=192.0.2.1
                           gateway = 192.0.2.1 (uncommitted)
T1:configuration net route (uncommitted)> commit
T1:configuration net routing> list
ROUTE      DESTINATION      GATEWAY    INTERFACE  TYPE    STATUS
...
route-003  192.0.2.101/32   192.0.2.1  ixgbe1     static  active
```

4. Create a replication target on S1.

The following example creates the replication target object on S1 to be used to replicate project Red from sp1 to tp1:

```
S1:shares replication targets>target
S1:shares replication target (uncommitted)> set hostname=192.0.2.103
                          hostname = 192.0.2.103 (uncommitted)
S1:shares replication target (uncommitted)> set label=t1-1
                             label = t1-1 (uncommitted)
S1:shares replication target (uncommitted)> set root_password=(set)
                     root_password = (set) (uncommitted)
S1:shares replication target (uncommitted)> commit
```

5. Create a replication action for each project.
   - Replicate project Red from pool sp1 to tp1
   - Replicate project Blue from pool sp1 to pool tp2
   - Replicate project Green from pool sp2 to tp2

   The following example creates the replication action for project Red:

   ```
   S1:> shares select Red replication action
   S1:shares Red action (uncommitted)> set target=t1-1
                               target=t1-1 (uncommitted)
   S1:shares Red action (uncommitted)> set pool=tp1
                                 pool=tp1 (uncommitted)
   S1:shares Red action (uncommitted)> commit
   ```

6. Set up to replicate project Blue from pool sp1 to tp2.

   Start with interface and address selection, and select interfaces S1/ixgbe3 and T2/ixgbe4, knowing that both S and T are in the CLUSTERED state and that the interface addresses are on the same subnet, 192.0.2.128/25. Next, define static routes on both appliances similar to the earlier examples. Then create replication target object t2-2 on S1, and create the replication action on S1 for project Blue using target object t2-2.

7. Set up to replicate project Green from pool sp2 to tp2.

   Start with interface selection and select interfaces S2/ixgbe2 and T2/ixgbe2. Create static routes on S2 and T2 using the selected interfaces and their addresses, define replication target object t2-1 using address of T2/ixgbe2, and finally create the replication action for project Green using target object t2-1.

8. Initiate replication for all three actions.
   a. Start with project Red:

      ```
      S1:> shares select Red replication select action-000
      S1:shares Red action-000> sendupdate
      ```

   b. Initiate replication for the actions for projects Blue and Green by following the previous example.

Replication Data Path Illustrated Examples

The following figures illustrate the replication data paths during replication updates for the replication actions for projects Red, Blue, and Green:

Normal Replication Data Path

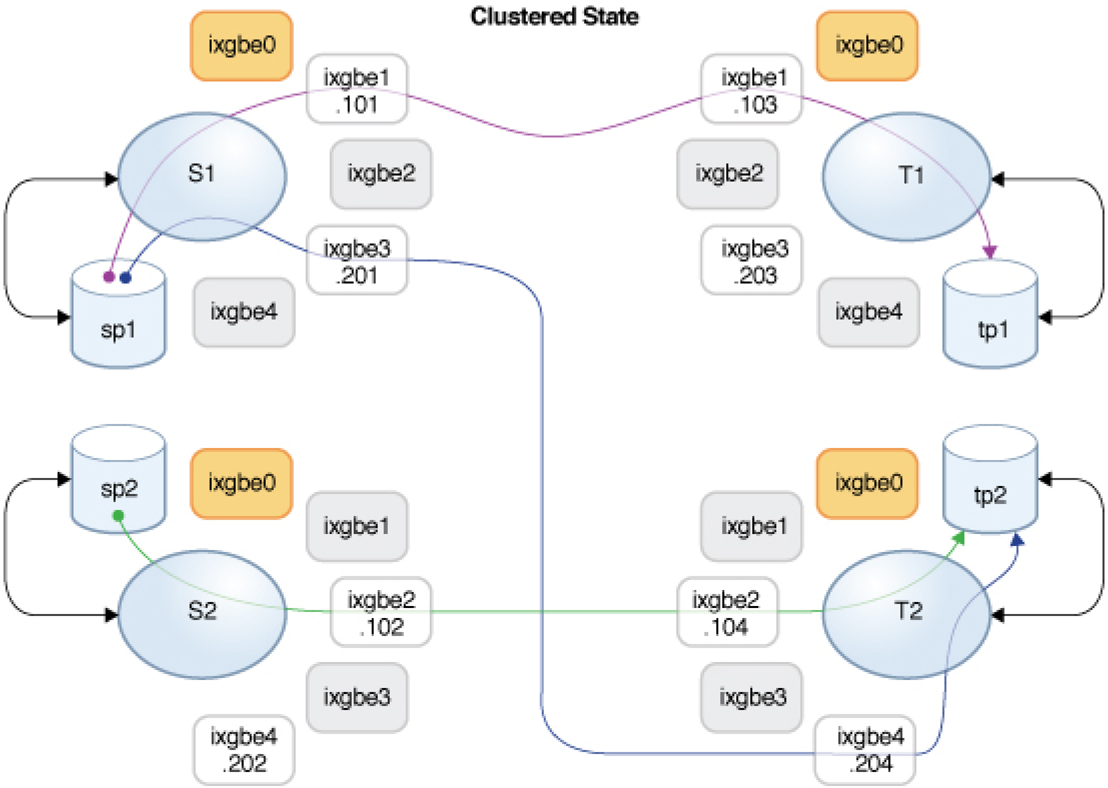

Assume that controller `T2` has been taken down for a maintenance. `T1` performed a takeover and now owns all of the resources. If replication updates for projects `Blue` and `Green` are in progress during the takeover, they will be canceled. After `T1` takes over, these replication updates can be resumed manually, or they will be resumed automatically if schedules are configured for the corresponding replication actions.

After controller `T1` has completed the takeover, it owns interfaces `ixgbe2` and `ixgbe4` which are necessary to continue replication updates for projects `Blue` and `Green`. The following figure shows the replication data path after `T1` completed the takeover.

Replication Data Path After T1 Takeover

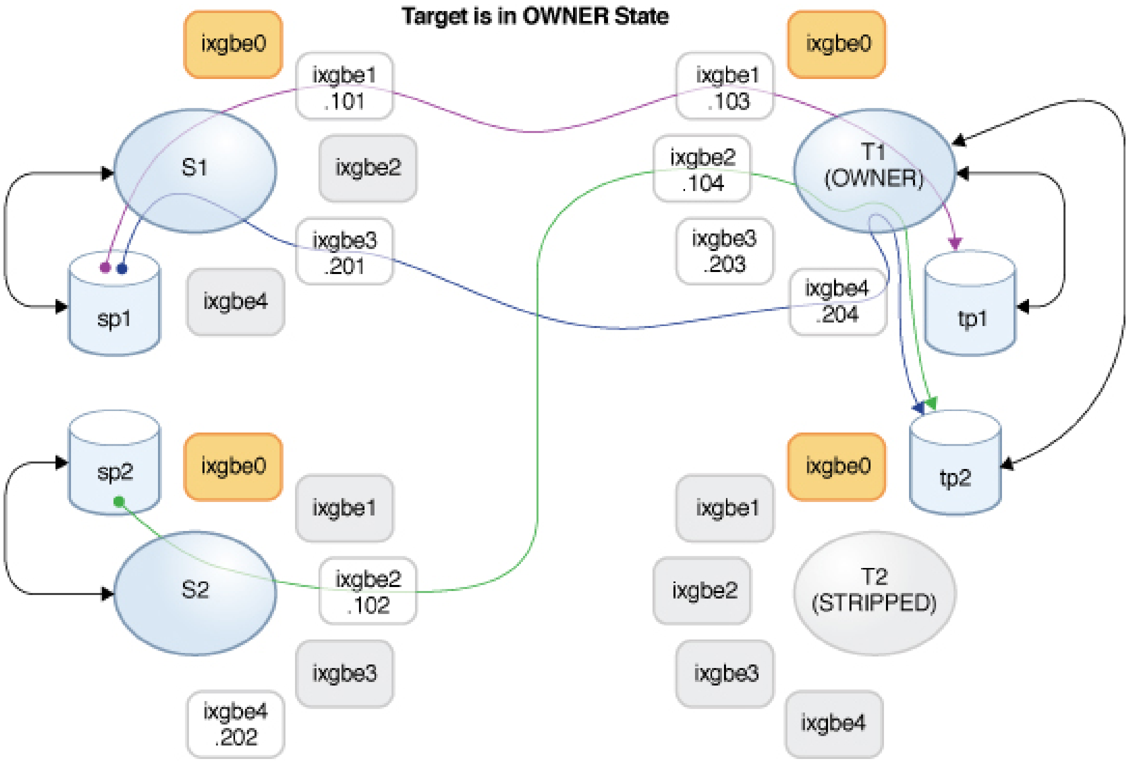

After `T2` is back online, a failback is performed on the `T1` controller and it takes over its resources. If replication updates of projects `Blue` and `Green` are in progress, they will be canceled and can be resumed following the completion of the failback.

Then controller `S2` is taken down for maintenance, and the takeover performed on the `S1` controller causes it to take ownership of all of the resources, including the interface required for continuing replication of project `Green`. If a replication update of project `Green` is in progress, it will be canceled and can be resumed following the completion of the takeover.

Data Path After Failback on T1 and Takeover on S1

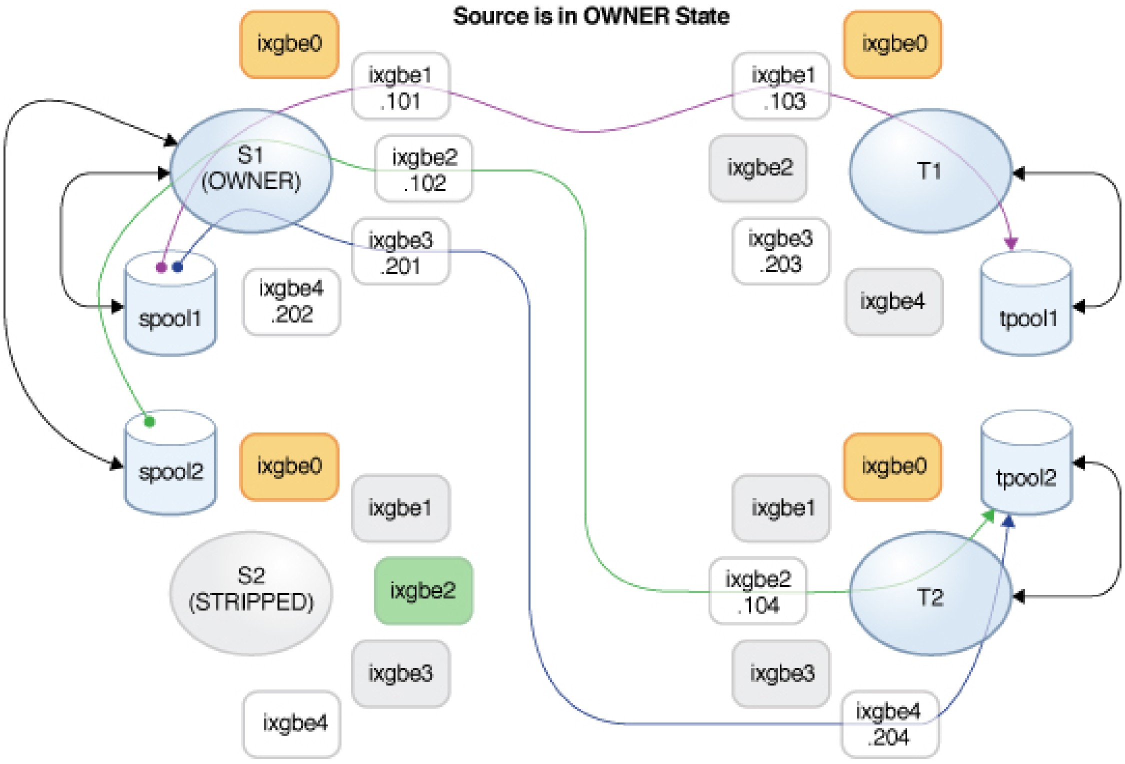

Related Topics

- Configuring Replication for a Clustered Configuration
- Remote Replication Workflow
- Remote Replication

Replication Snapshots and Data Consistency

The source appliance replicates snapshots atomically to the target, meaning the contents of the replica always exactly matches the source's share at the time the snapshot was taken. Because the snapshots for all shares sent in a particular group are taken at the same time, the entire package contents after the completion of a successful replication update exactly matches the group's content when the snapshot was created on the source.

However, each share's snapshots are replicated separately, so it is possible for some shares within a package to have been updated with a snapshot that is more recent than those of other shares in the same package. This is true during a replication update and after a failed replication update.

To summarize:

- Each share is always point-in-time consistent on the target.
- When no replication update is in progress and the previous replication update succeeded, each package's shares are also point-in-time consistent with each other.
- When a replication update is in progress or the previous update failed, package shares may be inconsistent with each other, but each one will still be self-consistent. If package consistency is important for an application, one must clone the replication package, which always clones the most recent successfully received snapshot of each share.

Related Topics

Replication Snapshot Management

Replication Snapshot Management

Snapshots are the basis for replication. The source and target must always share a common snapshot to continue replicating incrementally, and the source must know which is the most recent snapshot that the target has. To facilitate this, the replication subsystem creates and manages its own snapshots. Administrators generally do not need to be concerned with them, but the details are described here since snapshots can have significant effects on storage utilization.

Each replication update for a particular action consists of the following steps:

- Determine whether this is an incremental or full update based on whether:
  - An attempt was made to replicate this action before and
  - The target already has the necessary snapshot for an incremental update.
- Take a new project-level snapshot.
- Send the update. For a full update, send the entire group's contents up to the new snapshot. For an incremental update, send the difference between from the previous (base) snapshot and the new snapshot.
- Record the new snapshot as the base snapshot for the next update and destroy the previous base snapshot (for incremental updates). The base snapshot remains on the target until the next update is received at which point it is the first thing that is destroyed.

This has several consequences for snapshot management:

- During the first replication update and after the initial update when replication is not active, there is exactly one project-level snapshot for each action configured on the project or any share in the group. A replication action may create snapshots on shares that are in the same project as the share(s) in the group being replicated by the action, but that are not being sent as part of the update for the group.
- During subsequent replication updates of a particular action, there may be two project-level snapshots associated with the action. Both snapshots may remain after the update completes in the event of failure where the source was unable to determine whether the target successfully received the new snapshot (as in the case of a network outage during the update that causes a failure).
- None of the snapshots associated with a replication action can be destroyed by the administrator without breaking incremental replication. The system will not allow administrators to destroy snapshots on either the source or target that are necessary for incremental replication. To destroy such snapshots on the source, one must destroy the action (which destroys the snapshots associated with the action). To destroy such snapshots on the target, one must first sever the package (which destroys the ability to receive incremental updates to that package).
- Administrators must not rollback to snapshots created prior to any replication snapshots. Doing so will destroy the later replication snapshots and break incremental replication for any actions using those snapshots.
- Replication's usage of snapshots requires that administrators using replication understand space management on the appliance, particularly as it applies to snapshots.

Intermediate Replication Snapshots

A replication action can be set to include non-replication snapshots. When the property **Include Snapshots** is set, replication updates include the non-replication snapshots created after the previous replication update (or since the share's creation, in the case of the first full update). This includes automatic snapshots and administrator-created snapshots. This property can be disabled to skip these snapshots and send only the changes between replication snapshots with each update.

The action property `include_snaps` should be enabled to replicate any intermediate snapshots, including auto snapshots.

Related Topics

- Space Management for Shares
- Managing User-Generated Snapshots

Replication Automatic Snapshot Management

The automatic scheduled snapshots feature allows for automatically creating and destroying snapshots for projects and/or shares based on administrator-provided schedules. The schedule specifies when to create an automatic snapshot, and how many snapshots to retain. Several schedules can be created for a project or a share.

With remote replication, snapshots, including automatic snapshots, can be included in replication updates and will be available on the replication target as part of the corresponding replication package.

By default, automatic snapshot retention policies are synchronized between the source and the target, which means that the number of automatic snapshots retained on the target corresponds to the retention setting (`Keep At Most`) in the project's or share's snapshot schedule. Additionally, the target will have the same retention hold setting and number of snapshot schedules that cannot be deleted manually if the retention hold is locked.

By setting the automatic snapshot retention policies to the `independent` setting, replication actions can be configured to retain a separate, specific number of automatic snapshots on the target throughout replication updates. Also, the retention policy can be changed from off to locked, or locked to off. When setting the retention policy to locked, specify the number of snapshot schedules that cannot be deleted manually on the target. The number of locked snapshot schedules must be the same or the smaller than the `Keep At Most` setting. When all locked snapshots for this schedule on the target have exceeded the `Keep At Most` property value, the retention hold changes from locked to off.

To use the snapshot retention hold feature, apply deferred update "Support for Snapshot Retention." For information about deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The following user role authorizations are required to make changes to a retention policy:

- Modify an automatic snapshot schedule: `scheduleSnap`
- Modify the retention hold value from off to locked: `scheduleLockedSnap`
- Modify the retention hold value from locked to off: `releaseSnapRetention`

Reverse Replication and Automatic Snapshot Management

When reversing replication, automatic snapshot retention settings are preserved: The source and target will continue to maintain their retention settings.

Example:

- `Source A` has configured automatic snapshots and retains 5 snapshots on `Source A`.
- Through a replication action on `Source A`, `Target B` has been configured to retain 10 automatic snapshots.

After reverse replication, the source and target have switched to `Source B` and `Target A`.

- Now `Source B` has the automatic snapshot schedule, still retaining 10 snapshots.
- `Target A` is still configured to retain 5 snapshots. This retention setting is now configurable through the replication action on `Source B`.

Performing another reverse replication will revert the source and target to their original configurations. For more information on configuring automatic snapshot retention on a target, see:

- Configuring Automatic Snapshot Retention on a Target (BUI)
- Configuring Automatic Snapshot Retention on a Target (CLI)

iSCSI Configurations and Replication

Replication updates include most of the configuration specified on the **Shares** screen for a project and its shares. This includes any target groups and initiator groups associated with replicated LUNs.

When using non-default target groups and initiator groups, administrators must ensure that the target groups and initiator groups used by LUNs within the project also exist on the replication target. If the target group or initiator group does not exist on the target system, a clone, sever, or reverse replication will fail. An error message reports that the initiator or target group name was either deleted or renamed on the target system.

The SCSI GUID associated with a LUN is replicated with the LUN. As a result, the LUN on the target appliance will have the same SCSI GUID as the LUN on the source appliance. Clones of replicated LUNs, however, will have different GUIDs (just as clones of local LUNs have different GUIDs than their origins).

Related Topics

- Remote Replication Workflow
- Remote Replication

Resumable Replication

When a scheduled or continuous replication update is interrupted due to a network failure, system outage, or operator action, data transfer automatically resumes from the point of interruption if enough data (approximately 25M) was already sent. This feature is available with software release OS8.7.0 or later, which must be installed both on the source and on the replication target. The estimated data size, shown in the replication progress monitor, includes data to be sent as part of both the original replication update and the resumed replication update.

If a failure occurs during a manual replication update, the update is not automatically retried. However, the data transfer will resume from the point of interruption with the next replication update.

Replication Alerts

Oracle ZFS Storage Appliance posts alerts when any of the following replication events occurs:

- A manual or scheduled replication update starts or finishes successfully (both source and target). See Start and Finish Alerts.
- An administrator explicitly cancels an update (both source and target).
- An update fails (both source and target).
- A scheduled update is skipped because another update for the same action is already in progress.
- A continuous replication starts for the first time, fails, or resumes after a failure.
- A replica time lag exceeds its specified threshold.

Start and Finish Alerts

Depending on the number of projects that are replicating and the frequency of the replication schedule, the number of start and finish alerts for scheduled updates can obscure other important alerts. To disable start and finish alerts for scheduled updates, see Setting Replication Start and Finish Alerts - BUI, CLI.

If replication start and finish alerts are enabled, then when replication to a target completes successfully, the system displays the finish alert at the top of the BUI window, in addition to recording the alert in the alert log.

The following figure shows a replication update finish alert at the top of a BUI window.

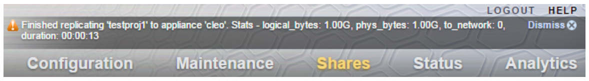

Finished replicating 'testproj1' to appliance 'cleo'. Stats - logical_bytes: 1.00G, phys_bytes: 1.00G, to_network: 0, duration: 00:00:13

The following shows a sample replication update finish alert in the BUI **Maintenance: Logs: Alerts**, or in the CLI `maintenance logs`, then `select alert`, and then `show`.

```
entry-71160 2020-11-18 00:03:23 019a9e92-6b12-400b-aea7-b900a58d30a6, Finished
replicating 'proj-9' to appliance 'zfs-storage'. Action -
4b9b1450-4300-4892-b89d-e170832ef531.. Stats - logical_bytes: 225K, phys_bytes:
67.4K, to_network: 33.4K, duration: 00:00:57, Minor Alert
```

The following table describes the finish alert properties for each replication update. For statistics for the most recently completed update or for accumulated totals over the lifetime of this replication action, see table "Replication Action stats Node Properties (CLI Read-Only)" in Replication Action Properties. If start and finish alerts are disabled, you cannot view statistics for specific previous updates, but you can view the statistics for the most recently completed update. For more information about properties with `dedup` and `dd_` in their names, see Deduplicated Replication.

Table 7-7 Replication Update Finish Alert Properties

CLI Property	Description
`logical_bytes`	Number of bytes that the replication update data stream would have contained if the data on disk had not been compressed and without any subsequent compression or deduplication.

Table 7-7 (Cont.) Replication Update Finish Alert Properties

CLI Property	Description
`phys_bytes`	Number of bytes in the internal replication data stream prior to replication deduplication or replication data stream compression.
`to_network`	Number of bytes that the replication data stream compression pipeline delivered to the network. This shows the consequence of replication data stream compression, if enabled.
`after_dedup`	Number of bytes in the internal replication data stream after deduplication of the replication data stream.
`duration`	Elapsed time to perform the replication update.
`dd_table_build`	Time to build the deduplication tables prior to the transmission of the replication update.
`dd_table_mem`	Maximum amount of memory that was consumed by the deduplication tables.

Replication Failures

Individual replication updates can fail for a number of reasons. Oracle ZFS Storage Appliance reports the reason for the failure in alerts posted on the source appliance or replication target, or on the **Replication** screen for the action that failed. You may be able to get details on the failure by clicking the orange alert icon representing the action's status.

The following are some common replication failures:

Failure	Details
Cancelled	The replication update was cancelled by an administrator. Replication can be cancelled on the source or target.
Network connectivity failure	The appliance was unable to connect to the replication target due to a network problem. Check for a misconfiguration on the source, target, or the network.
Peer verification failed	The appliance failed to verify the identity of the target. This occurs most commonly when the target has been reinstalled or factory reset. A new replication target must be configured on the source appliance for a target which has been reinstalled or factory reset in order to generate a new set of authentication keys. See Replication Targets.
Peer RPC failed	A remote procedure call failed on the target system. This occurs most commonly when the replication target is running incompatible software. See Oracle ZFS Storage Appliance: Remote Replication Compatibility [Doc ID 1958039.1] at https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1 .
Name collision	Replication of <*project/share*> from <*source*> failed due to a name collision with `@`*`<snapname>`* being held on the target for NDMP. To recover, rename (or remove) the snapshot on the replication source that has the same name as the snapshot held by NDMP on the target (the one named in the alert), unless it starts with `.rr`. Then either perform a manual sync or allow the replication source to automatically retry the replication update.

Failure	Details
No package	Replication failed because no package exists on the target to contain the replicated data. Since the package is created when configuring the action, this error typically happens after an administrator has destroyed the package on the target. This error could also occur if the storage pool containing the package is not imported on the target system, which may occur if the pool is faulted or if storage or networking has been reconfigured on the replication target.
Disabled	Replication failed because it is disabled on the target. Either the replication service is disabled on the target or replication has been disabled for the specific package being replicated.
Target busy	Replication failed because the target system has reached the maximum number of concurrent replication updates. The system limits the maximum number of ongoing replication operations to avoid resource exhaustion. When this limit is reached, subsequent attempts to receive updates will fail with this error, while subsequent attempts to send updates will queue up until resources are available.
Target is missing	The most recent replication update failed because the target is missing. If the target is no longer configured on the source, the action will be permanently disabled. If this error occurs, destroy the replication action and reconfigure the replication target and action.
Out of space	Replication failed because the source system had insufficient space to create a new snapshot. This may be because there is no physical space available in the storage pool or because the project or one of its shares would be over quota because of reservations that do not include snapshots.
Key Unavailability	Replication failed because the encryption key used by the share is not available either on the source or target system. Review the alerts on both the source and replication target to ensure the key is available on both systems. See Backing Up, Replicating, and Restoring Encrypted Projects and Shares for information about replicating encrypted shares and projects.
Incompatible target	Replication failed because the target system is unable to receive the source system's data stream format. This can happen as a result of upgrading a source system and applying deferred updates without having upgraded and applied the same updates on the target. For deferred updates that have remote replication implications, see Oracle ZFS Storage Appliance: Remote Replication Compatibility [Doc ID 1958039.1] at https://support.oracle.com/epmos/faces/DocumentDisplay?id=1958039.1 .
iSCSI initiator/target missing	A replication clone, sever, or reverse operation failed because the initiator group or target group LUNs do not exist for the LUNs included in the replication package. The initiator or target group name was either deleted or renamed on the replication target.
Misc	Replication failed, but no additional information is available on the source. Check the alert log on the target system and if necessary contact support for assistance. Some failure modes that currently fall into this category include insufficient disk space on the target to receive the update and attempting to replicate a clone whose origin snapshot does not exist on the target system.

A replication update fails if any part of the update fails. The shares inside a project are replicated serially and changes are not rolled back from a failed update. As a result, when an update fails, some shares on the target may be up to date while others are not. For more information, see Replication Snapshots and Data Consistency.

When a scheduled or continuous replication fails, the system waits several minutes and tries again. The system will continue retrying failed scheduled or continuous replications indefinitely. At any point during the retry procedure, initiating a manual update will immediately begin a retry, circumventing the usual delay between successive retries. If the manual update completes successfully, it terminates the retry sequence and the replication action reverts to its normal scheduled or continuous updates.

For more information about failed or interrupted replication updates, see Resumable Replication.

When a replication update is in progress and another update is scheduled, the scheduled replication is deferred until the previous update completes, and an alert is posted.

Related Topics

How to Troubleshoot Replication Issues (Doc ID 1397959.1) on My Oracle Support (https://support.oracle.com/)

Compressed Replication

The compressed replication feature improves performance when replicating compressible data between source and target sites that have limited bandwidth. Before a replication stream is sent to the target, it is automatically compressed at a rate based on current CPU utilization and network I/O throughput. The replication stream is then decompressed when received by the replication target. If any part of the data is not compressible, that portion will be sent as if compression were disabled.

All replication streams will be compressed, unless you explicitly disable compression. If your WAN equipment provides compression, for example through a WAN accelerator, disable the compression feature by following the procedure Disabling Replication Compression - BUI, CLI.

The source appliance and replication target require software version 2013.1.4.0 or later to support replication compression. If the target has an earlier version, a warning icon ● is displayed next to the target name. You will need to update the replication target to at least the minimum version.

You can view replication performance statistics in the BUI on the source appliance, under the progress bar for the replication.

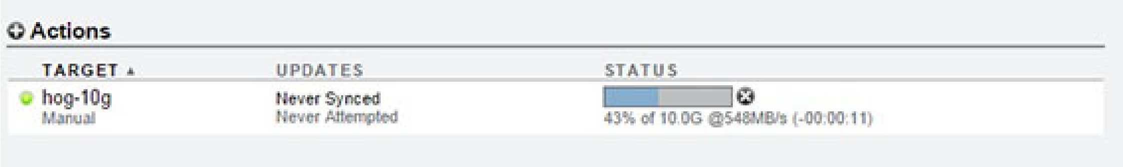

TARGET ▲	UPDATES	STATUS
● hog-10g Manual	Never Synced Never Attempted	43% of 10.0G @548MB/s (-00:00:11)

Replication Packages

Packages are containers for replicated projects and shares. Each replication action on a source appliance corresponds to one package on a replication target.

You can browse replicated projects, shares, snapshots, and properties much like local projects and shares, using the BUI or CLI. However, because replicated shares must exactly match their counterparts on the source appliance, many management operations are not allowed inside replication packages.

You can modify the following properties of replicated projects and shares:

- **Reservation, Compression, Copies, Deduplication, and Caching** - These properties can be changed on the replication target to effect different cost, flexibility, performance, or reliability policies on the replication target from the source.
- **Mountpoint and Sharing Properties** (for example, `sharenfs`, `SMB resource name`) - These properties control how shares are exported to NAS clients and can be changed to effect different security or protection policies on the replication target from the source.

Such property modifications persist across replication updates.

Managing Replication Package Properties

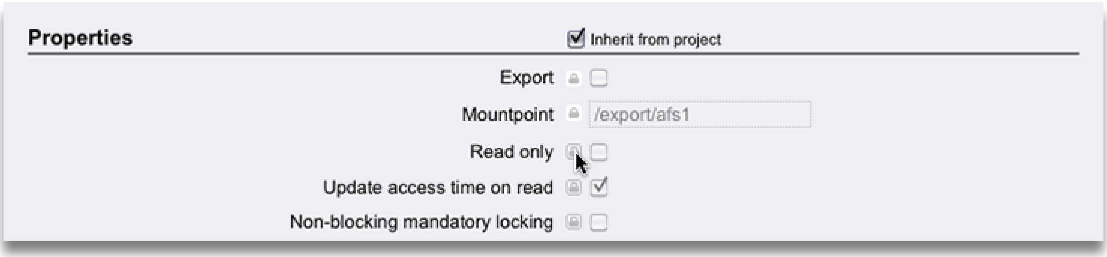

Properties ☑ Inherit from project

Export ☐

Mountpoint /export/afs1

Read only ☐

Update access time on read ☑

Non-blocking mandatory locking ☐

Related Topics

- Project and Share Properties
- Severing Replication

Cloning a Replication Package or Share

A *clone* of a replicated package is a local, mutable project that can be managed like any other project on Oracle ZFS Storage Appliance. When the clone project is created, the most recently received snapshot of the replicated shares is used to create the shares within the clone project. These clones share storage with their origin snapshots the same way that clones of share snapshots do (see Cloning a Snapshot - BUI, CLI). This mechanism can be used to failover in the case of a catastrophic problem at the replication source, or simply to provide a local version of the data that can be modified.

As long as a clone exists, its origin snapshot cannot be destroyed. When destroying the snapshot (possibly as a result of destroying the share, project, or replication package of which the snapshot is a member), the system warns administrators of any dependent clones that will be destroyed by the operation. Note that snapshots can also be destroyed on the source at any time and such snapshots are destroyed on the target as part of the subsequent replication update. If such a snapshot has clones, the snapshot will not be destroyed until the last clone has been destroyed.

Replicating Clones

When replicating clones, it is important to understand the relationship between a clone replica and its origin snapshot. By default, the replica of a clone maintains its relationship with its origin snapshot, mandating that a replica of the origin snapshot also exist on the target. A replica of a clone origin snapshot must reside in the same pool as the clone, but does not have to be in the same project.

To maintain the relationship between a replicated clone and its origin snapshot, the origin snapshot must be:

- Replicated to the target before the initial replication of the clone or
- Replicated as part of the same update.

This restriction is not enforced by the appliance software, but must be followed to ensure a successful replication update.

There are several ways to ensure successful replication of a clone so it maintains its relationship with its origin snapshot:

- If the clone's origin snapshot is in the same project, use project-level replication.
- If the share containing the clone origin snapshot is not in the same project or if the clone or its origin share have been omitted from project-level replication, replicate the origin share first and then replicate the clone using project-level or share-level replication.
- On the target system, do not destroy the origin of the clone unless you also intend to destroy the clone itself.

To ensure that the origin snapshot is sent to the target, always set the property `Include snapshots` for the origin's replication action.

Just as a clone and its origin snapshot conserve space on the source appliance, a replicated clone and its replicated origin snapshot conserve space on the replication target. If space conservation on the replication target is less important, the administrator may set the property `Include clone origin as data`. When this property is set, and the origin snapshot of a clone is *not* replicated in the same update as the clone, the source appliance inserts a copy of the clone origin's data content into the replica clone. Thus, there is no need to replicate the clone origin share first, but the copy of the clone origin data consumes additional storage space on the target.

When `Include snapshots` and `Include clone origin as data` are both set, the replica clone contains only the snapshots that are present in the clone on the source. The source appliance inserts the clone origin data content, not the clone origin snapshots, into the replica clone. This ensures that the snapshots present in the replica clone match the snapshots present in the clone on the source.

The property `Include clone origin as data` does not affect the replication of a clone and its origin snapshot when they are both replicated in the same update. When replicated together by the same replication action, the relationship between the clone and its origin snapshot is preserved and the space sharing benefit is retained on the target.

Related Topics

- Project vs. Share Replication
- Cloning a Replication Package - BUI, CLI
- Project and Share Properties

Multi-target Reversal

In a disaster recovery setup consisting of a source replicating a project to multiple targets, when a reverse is performed at one of the targets, the multi-target reversal feature provides the ability to continue sending incremental updates from the new source to all other originally configured targets.

Note:

Replication reversal is not supported to a filesystem with mandatory file retention. For information about the file retention feature, see File Retention Management.

Basic Setup

A typical multi-target reversal setup consists of the following entities:

- **Source** - A replication source responsible for sending out replication updates and performing snapshot management. There should be exactly one source.
- **Potential source** - A potential source is a replication target in the multi-target reversal configuration that needs to take over the job of sending out incremental replication updates (along with other tasks such as snapshot management, and so on) to all the targets. There should at least one potential source.
- **Dedicated target** - A dedicated target is a replication target in the multi-target reversal configuration that cannot perform a reverse operation. The dedicated target needs fewer snapshots than a potential source. An appliance running an older version of software can be configured as a dedicated target in a multi-target reversal setup. Dedicated targets may or may not be present in a multi-target reversal setup.
- **Multi-target reversal group** - A group of replication actions for the same dataset that allows performing a reverse on one of the potential sources. Following the reverse, the new source can continue sending incremental updates to all originally configured targets.

In the following figure, this example of a multi-target reversal setup consists of a source (`S`), three potential sources (`PS1`, `PS2`, `PS3`), and a dedicated target (`DT1`). The source is responsible to send updates to the targets. When the source goes down, one of the potential sources, in this case `PS1`, is chosen to be the new source.

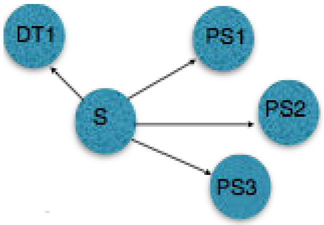

When a reverse is performed on this new source, it continues to send incremental updates to the other targets, as shown in the following figure.

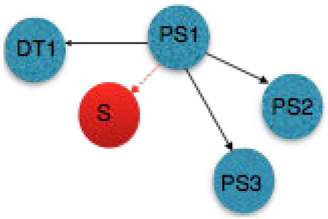

Multi-target Reversal Group Management

A multi-target group is created by configuring at least one project-level action as a potential source. In the absence of a potential source, the multi-target reversal group is not created, and reversal on any target is permitted, but the reverse will only create a replication action to replicate back to the original source.

> **Note:**
>
> While using multi-target reversal, multiple replication snapshots may be maintained for every action in the multi-target reversal group.

Creating a Multi-target Reversal Group

1. Configure a potential source.

   The potential source can be configured on a project-level replication action as shown in the following BUI figure. Configuration changes to a replication action are propagated as part of a replication update. There are no guarantees that such configuration changes will be delivered to all the members of the multi-target group at the same time.

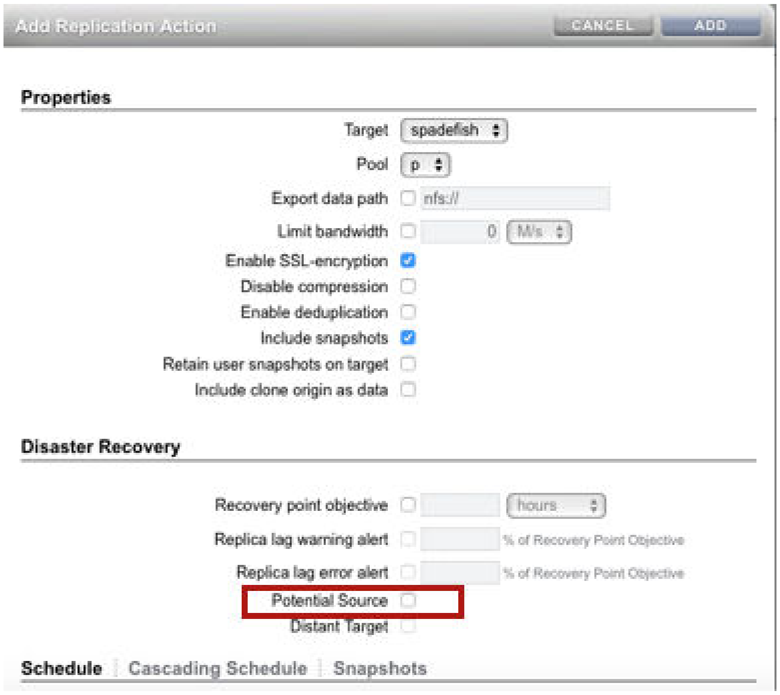

Properties

Target: spadefish
Pool: p
Export data path □ nfs://
Limit bandwidth □ 0 M/s
Enable SSL-encryption ☑
Disable compression □
Enable deduplication □
Include snapshots ☑
Retain user snapshots on target □
Include clone origin as data □

Disaster Recovery

Recovery point objective □ hours
Replica lag warning alert □ % of Recovery Point Objective
Replica lag error alert □ % of Recovery Point Objective
Potential Source □
Distant Target □

2. Create replication targets on the potential sources as described in Creating a Replication Target (BUI) or Creating a Replication Target (CLI).

   When a reverse is performed on a potential source, actions to all original targets will be created. If there are no replication targets configured on the potential source, the newly created actions will not have replication targets. Therefore, these new actions will be unbound. When appropriate replication targets are created at later time, the unbound actions will be automatically bound to correct targets.

> **Note:**
>
> Because the IP address of the target is used by the auto-bind process, targets to the same appliances and the same IP addresses as used on the original source need to be created.

Viewing Unbound Actions on the Source using the CLI

a. Select the appropriate project.

b. Go under the `replication` sub-node of that project.

c. Use the `show` command to display actions as shown in the following figure. Unbound actions have an `<undefined>` target.

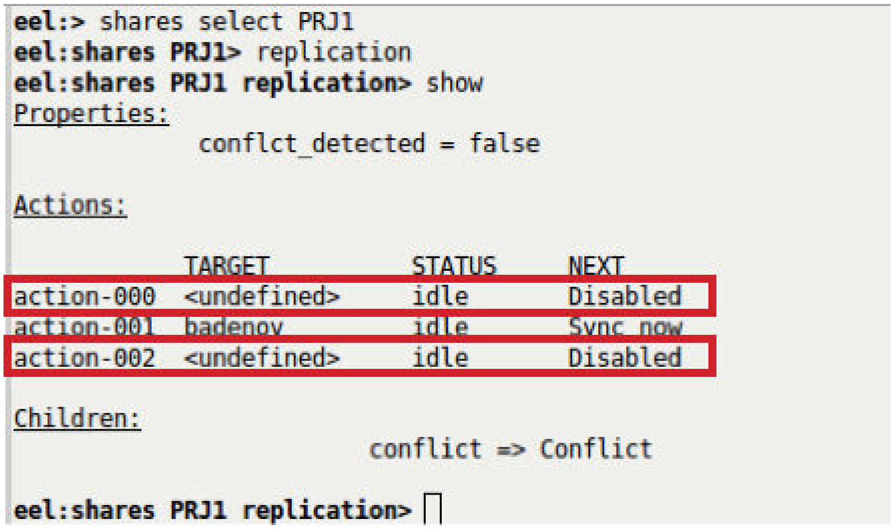

```
eel:> shares select PRJ1
eel:shares PRJ1> replication
eel:shares PRJ1 replication> show
Properties:
                    conflct_detected = false

Actions:

                  TARGET             STATUS      NEXT
action-000   <undefined>        idle        Disabled
action-001   badenov            idle        Sync now
action-002   <undefined>        idle        Disabled

Children:
                                   conflict => Conflict

eel:shares PRJ1 replication> 
```

Viewing Unbound Actions on the Source using the BUI

a. Select the appropriate project.

b. Click **Replication**.

c. Unbound actions have an undefined target (**<None>**).

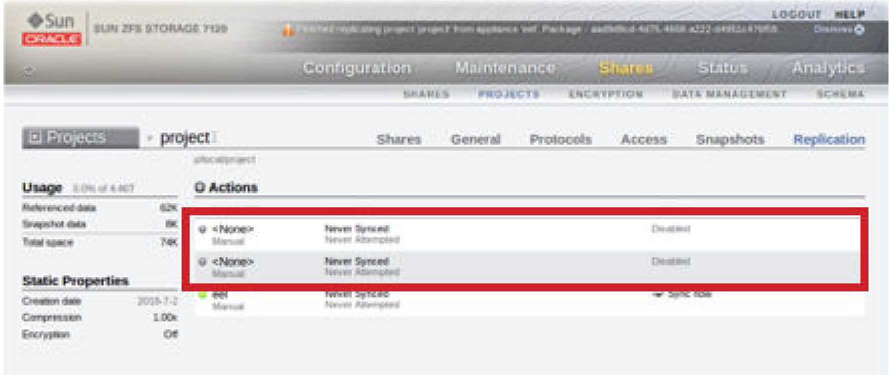

Referenced data	62K
Snapshot data	8K
Total space	74K
Creation date	2016-7-2
Compression	1.00x
Encryption	Off

<None> Manual	Never Synced Never Attempted	Disabled
<None> Manual	Never Synced Never Attempted	Disabled
eel Manual	Never Synced Never Attempted	Sync now

Monitoring Potential Targets

Replication packages on potential sources show a list of potential targets. During the reverse actions to potential targets will be created in addition to the action to replicate back to the original source. The list includes the target name, or IP address, and the package ID on the corresponding target.

Potential targets can be monitored as shown in the following CLI figure.

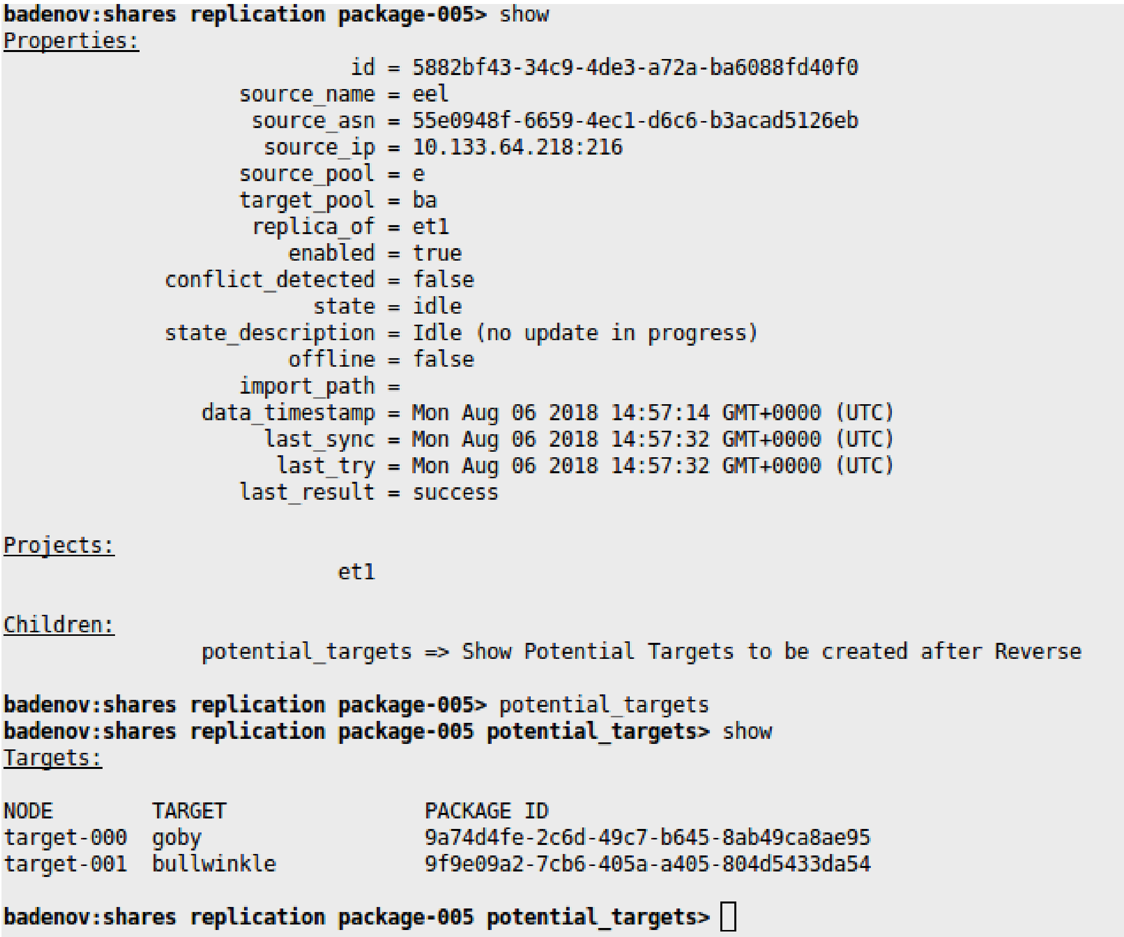

```
badenov:shares replication package-005> show
Properties:
                            id = 5882bf43-34c9-4de3-a72a-ba6088fd40f0
                   source_name = eel
                    source_asn = 55e0948f-6659-4ec1-d6c6-b3acad5126eb
                     source_ip = 10.133.64.218:216
                   source_pool = e
                   target_pool = ba
                    replica_of = et1
                       enabled = true
             conflict_detected = false
                         state = idle
             state_description = Idle (no update in progress)
                       offline = false
                   import_path =
                data_timestamp = Mon Aug 06 2018 14:57:14 GMT+0000 (UTC)
                     last_sync = Mon Aug 06 2018 14:57:32 GMT+0000 (UTC)
                      last_try = Mon Aug 06 2018 14:57:32 GMT+0000 (UTC)
                   last_result = success

Projects:
                           et1

Children:
                potential_targets => Show Potential Targets to be created after Reverse

badenov:shares replication package-005> potential_targets
badenov:shares replication package-005 potential_targets> show
Targets:

NODE        TARGET                PACKAGE ID
target-000  goby                  9a74d4fe-2c6d-49c7-b645-8ab49ca8ae95
target-001  bullwinkle            9f9e09a2-7cb6-405a-a405-804d5433da54

badenov:shares replication package-005 potential_targets>
```

Potential targets can be monitored as shown in the following BUI figure.

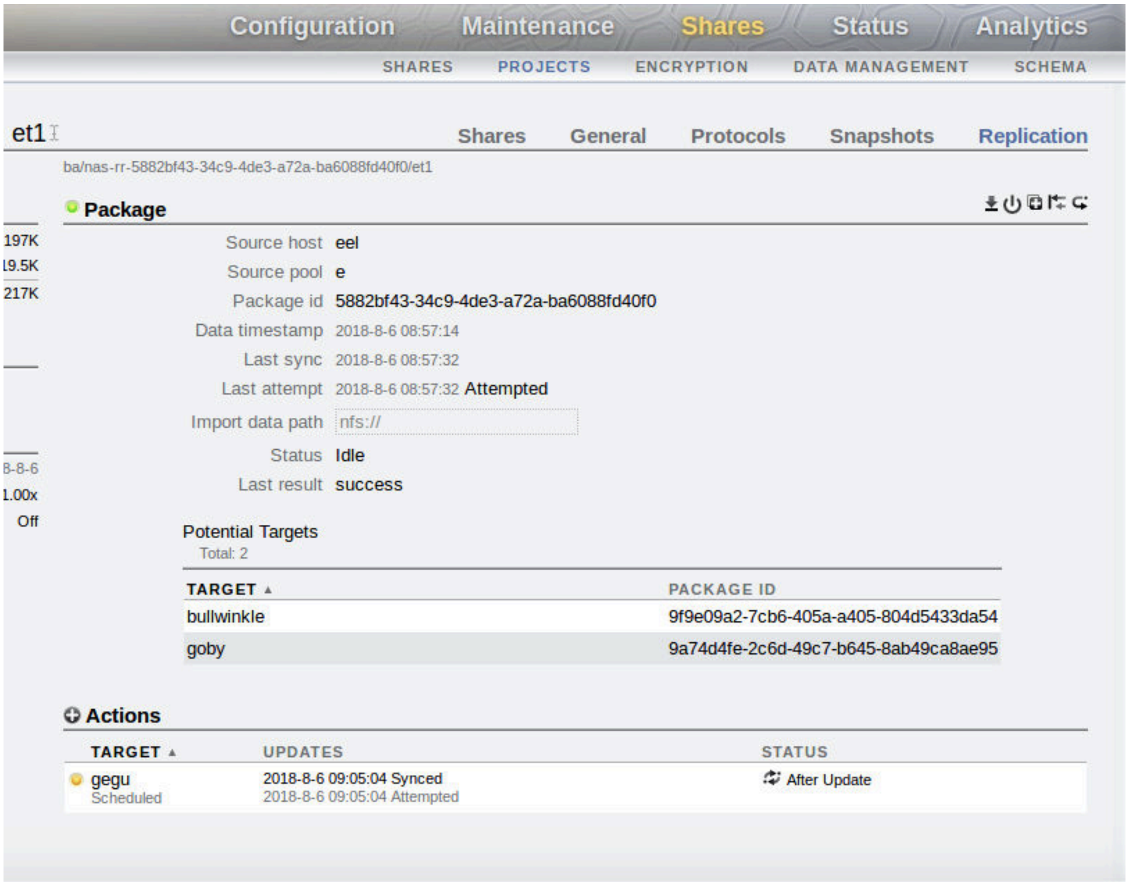

et1

ba/nas-rr-5882bf43-34c9-4de3-a72a-ba6088fd40f0/et1

Package

Source host	eel
Source pool	e
Package id	5882bf43-34c9-4de3-a72a-ba6088fd40f0
Data timestamp	2018-8-6 08:57:14
Last sync	2018-8-6 08:57:32
Last attempt	2018-8-6 08:57:32 Attempted
Import data path	nfs://
Status	Idle
Last result	success

Potential Targets
Total: 2

TARGET ▲	PACKAGE ID
bullwinkle	9f9e09a2-7cb6-405a-a405-804d5433da54
goby	9a74d4fe-2c6d-49c7-b645-8ab49ca8ae95

Actions

TARGET ▲	UPDATES	STATUS
gegu Scheduled	2018-8-6 09:05:04 Synced 2018-8-6 09:05:04 Attempted	After Update

Conflict Detection and Resolution

In the following figure of a multi-target reversal group, A is the source of the group.

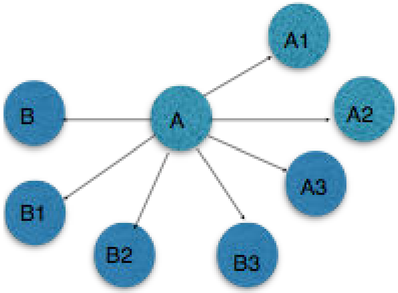

At any point of time, only one source should exist in a multi-target reversal group. A conflict can arise when more than one source may exist, as in the following scenarios.

- Due to an administration error when a reverse is performed on two or more potential sources receiving updates from the same replication source.
- In the event of a network segmentation, the administrator may choose to create multiple multi-target reversal segments, each with its own source. After recovering from such a network segmentation, sources from each of the multi-target reversal segments would attempt to send updates to each other, and this results in a conflict.

As shown in the following figure, there are now two sources: A1 and B.

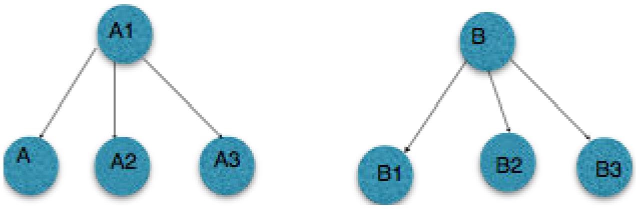

Detecting Conflict on the Source(s) and Target(s)

1. As shown in the following figure, conflict is detected on the targets when two or more sources attempt to send updates to the same target. In this scenario, updates from one source will succeed, and that source will be able to continue sending successful updates to the target. Updates from other sources will fail with a specific alert, and the target will show conflict notification prompting the administrator to resolve the conflict by selecting the correct source. After the conflict has been resolved, the target will receive the update from the source that was set by the administrator. However, if the incorrect source continues to send an update, a new alert will be raised at the incorrect source, and a conflict resolution notification will again be raised at the target.

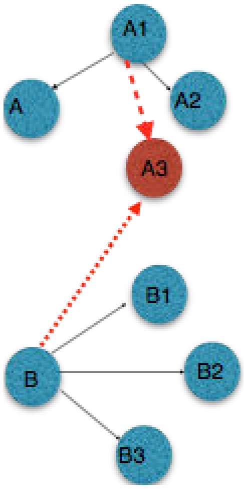

2. Conflict is detected when the source (node `B`) of one group sends an update to a source in the other group (node `A1`) as shown in the following figure. The update on the sending node (node `B`) will fail with an alert, and the source that receives the update (node `A1`) will show a conflict notification, prompting the administrator to resolve the conflict by selecting the correct source. After the conflict has been resolved, the incorrect source will convert its project to a package on the next update and, therefore, becomes a target.

> **Note:**
>
> When targets receive updates from incorrect sources, data rollback may be performed during the first update from the correct source.

Example of a CLI alert raised on the source that caused the conflict:

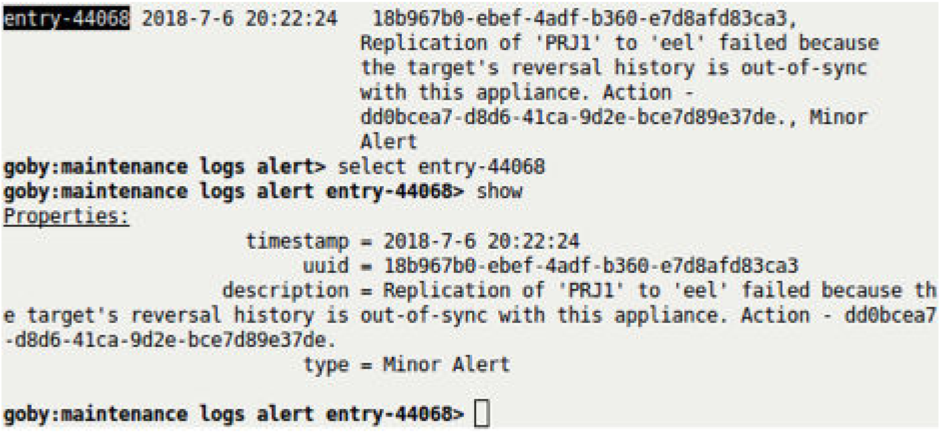

```
entry-44068 2018-7-6 20:22:24   18b967b0-ebef-4adf-b360-e7d8afd83ca3,
                               Replication of 'PRJ1' to 'eel' failed because
                               the target's reversal history is out-of-sync
                               with this appliance. Action -
                               dd0bcea7-d8d6-41ca-9d2e-bce7d89e37de., Minor
                               Alert
goby:maintenance logs alert> select entry-44068
goby:maintenance logs alert entry-44068> show
Properties:
                     timestamp = 2018-7-6 20:22:24
                          uuid = 18b967b0-ebef-4adf-b360-e7d8afd83ca3
                   description = Replication of 'PRJ1' to 'eel' failed because th
e target's reversal history is out-of-sync with this appliance. Action - dd0bcea7
-d8d6-41ca-9d2e-bce7d89e37de.
                          type = Minor Alert

goby:maintenance logs alert entry-44068> 
```

Example of a BUI alert raised on the source that caused the conflict:

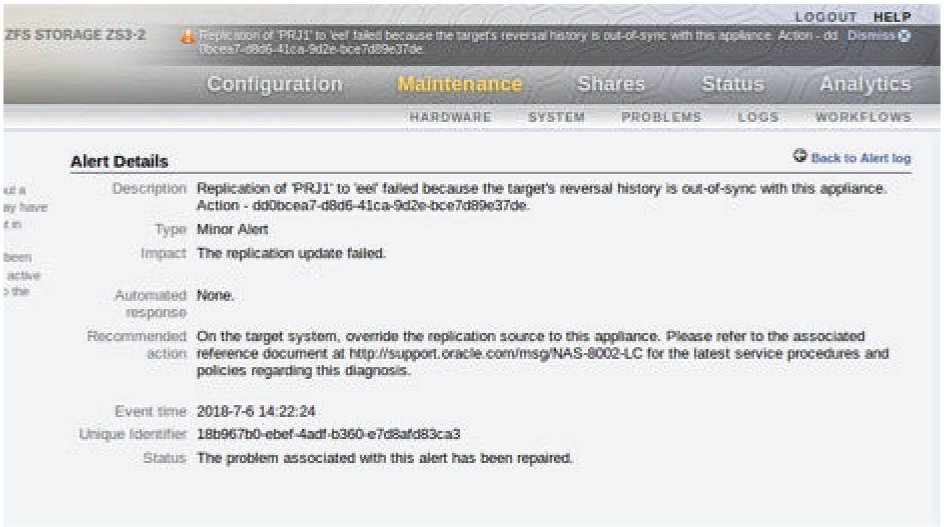

Replication of 'PRJ1' to 'eel' failed because the target's reversal history is out-of-sync with this appliance. Action - dd0bcea7-d8d6-41ca-9d2e-bce7d89e37de

Alert Details

Description: Replication of 'PRJ1' to 'eel' failed because the target's reversal history is out-of-sync with this appliance. Action - dd0bcea7-d8d6-41ca-9d2e-bce7d89e37de.

Type: Minor Alert

Impact: The replication update failed.

Automated response: None.

Recommended action: On the target system, override the replication source to this appliance. Please refer to the associated reference document at http://support.oracle.com/msg/NAS-8002-LC for the latest service procedures and policies regarding this diagnosis.

Event time: 2018-7-6 14:22:24

Unique Identifier: 18b967b0-ebef-4adf-b360-e7d8afd83ca3

Status: The problem associated with this alert has been repaired.

General Steps for Resolving a Conflict

1. Disable updates from all sources.
2. Select the correct source.
3. On the target(s), resolve the conflict by setting the correct source using the following CLI or BUI Target procedure.
4. On the incorrect source(s), resolve the conflict by setting the correct source using the following CLI or BUI Source procedure.

5. Enable updates from the correct source.

> **Note:**
>
> 1) Updates from other sources will cause corresponding conflict notifications to show up. 2) The conflict resolution tool does not store information persistently. If the node restarts for any reason, conflict and resolution selections will be discarded, and any updates from sources that had caused conflict before the restart will show the conflict notification again.

Target: Resolving the Conflict using the CLI

1. Select the appropriate replication package, go under the project, and then select `replication`.
2. Check the value of property `conflict_detected`. When it is `true`, it indicates that conflict has been detected.
3. Go under sub-node `conflict` and resolve the conflict by setting the value of property `source` to the correct source, as shown in the following figure.

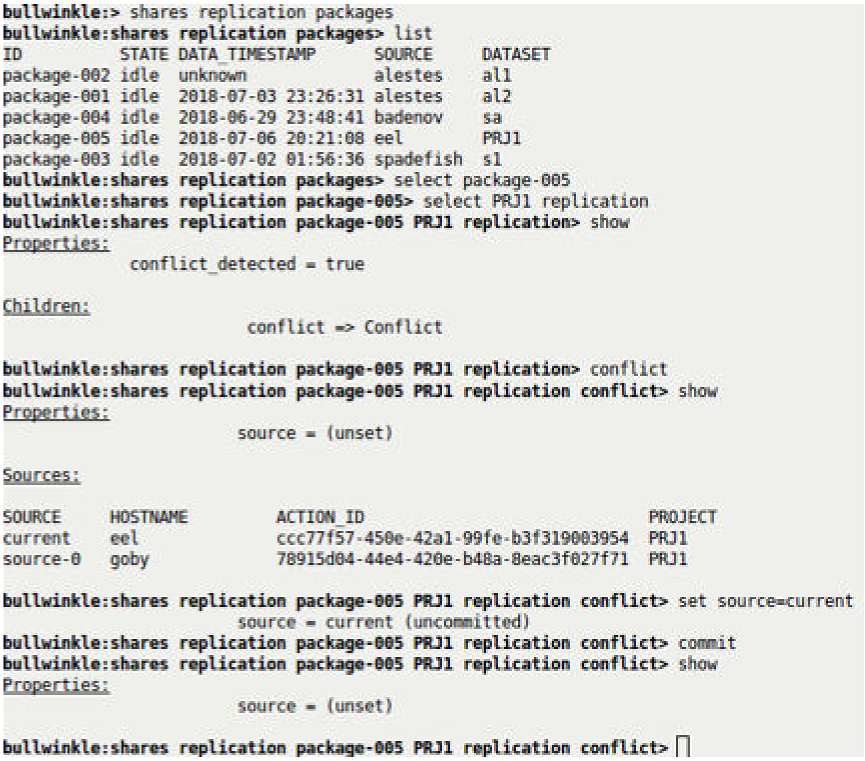

```
bullwinkle:> shares replication packages
bullwinkle:shares replication packages> list
ID          STATE DATA_TIMESTAMP      SOURCE     DATASET
package-002 idle  unknown             alestes    al1
package-001 idle  2018-07-03 23:26:31 alestes    al2
package-004 idle  2018-06-29 23:48:41 badenov    sa
package-005 idle  2018-07-06 20:21:08 eel        PRJ1
package-003 idle  2018-07-02 01:56:36 spadefish  s1
bullwinkle:shares replication packages> select package-005
bullwinkle:shares replication package-005> select PRJ1 replication
bullwinkle:shares replication package-005 PRJ1 replication> show
Properties:
             conflict_detected = true

Children:
                         conflict => Conflict

bullwinkle:shares replication package-005 PRJ1 replication> conflict
bullwinkle:shares replication package-005 PRJ1 replication conflict> show
Properties:
                        source = (unset)

Sources:

SOURCE     HOSTNAME         ACTION_ID                             PROJECT
current    eel              ccc77f57-450e-42a1-99fe-b3f319003954  PRJ1
source-0   goby             78915d04-44e4-420e-b48a-8eac3f027f71  PRJ1

bullwinkle:shares replication package-005 PRJ1 replication conflict> set source=current
                        source = current (uncommitted)
bullwinkle:shares replication package-005 PRJ1 replication conflict> commit
bullwinkle:shares replication package-005 PRJ1 replication conflict> show
Properties:
                        source = (unset)

bullwinkle:shares replication package-005 PRJ1 replication conflict> 
```

Target: Resolving the Conflict using the BUI

1. Select the replication package.
2. Click **Replication**.

   Conflict notification should be displayed, as shown in the following figure.

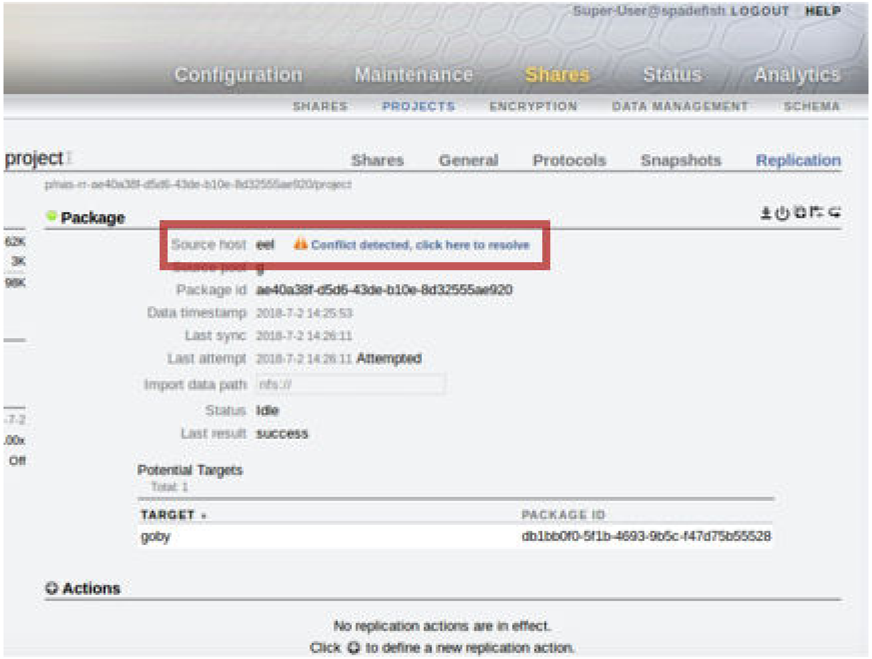

p/nas-rr-ae40a38f-d5d6-43de-b10e-8d32555ae920/project

Package

Source host eel Conflict detected, click here to resolve
Source pool g
Package id ae40a38f-d5d6-43de-b10e-8d32555ae920
Data timestamp 2018-7-2 14:25:53
Last sync 2018-7-2 14:26:11
Last attempt 2018-7-2 14:26:11 Attempted
Import data path nfs://
Status Idle
Last result success

Potential Targets
Total: 1

TARGET	PACKAGE ID
goby	db1bb0f0-5f1b-4693-9b5c-f47d75b55528

Actions

No replication actions are in effect.
Click ⊕ to define a new replication action.

3. Click on the conflict notification to open the replication conflict resolution dialog box, as shown in the following figure.

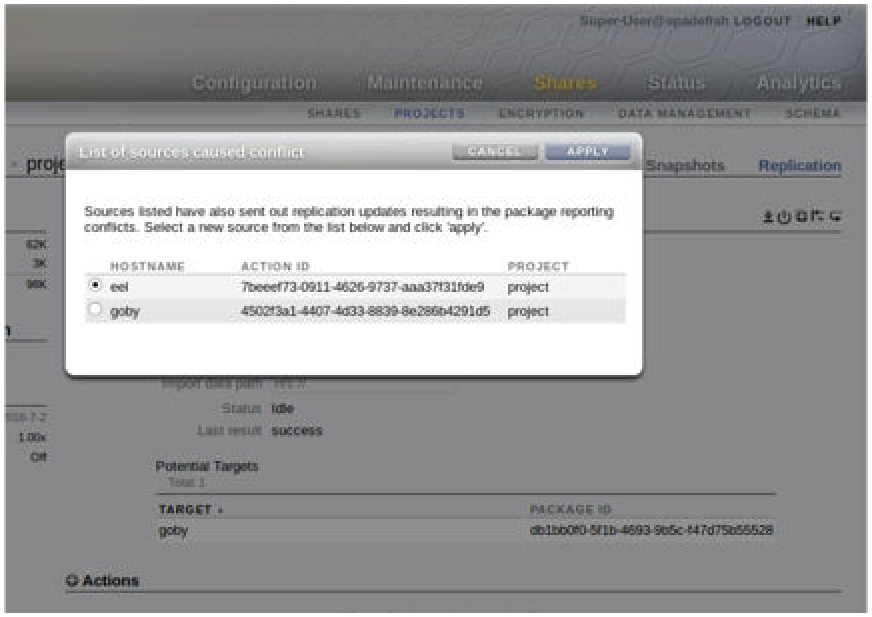

Sources listed have also sent out replication updates resulting in the package reporting conflicts. Select a new source from the list below and click 'apply'.

	HOSTNAME	ACTION ID	PROJECT
●	eel	7beeef73-0911-4626-9737-aaa37f31fde9	project
○	goby	4502f3a1-4407-4d33-8839-8e286b4291d5	project

Import data path nfs://
Status Idle
Last result success

Potential Targets
Total: 1

TARGET	PACKAGE ID
goby	db1bb0f0-5f1b-4693-9b5c-f47d75b55528

Actions

4. Select the correct source, and click **APPLY**.

   The updates from that source will now succeed.

Source: Resolving the Conflict using the CLI

1. Select the incorrectly reversed project.
2. Go under sub-node `replication` of that project.
3. Check the value of property `conflict_detected`. When it is `true`, it indicates that conflict has been detected.
4. Go under sub-node `conflict` and resolve the conflict by setting the value of property `source` to the correct source, as shown in the following figure.

   From this point, the update from the selected source will cause the project to be converted into a replication package.

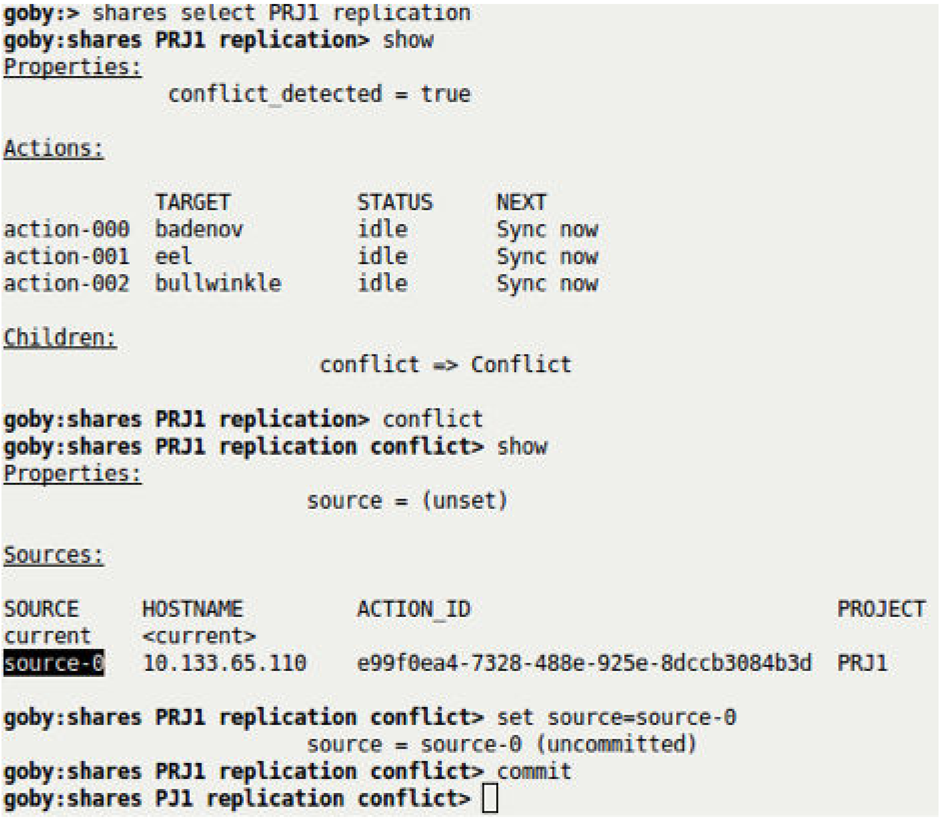

```
goby:> shares select PRJ1 replication
goby:shares PRJ1 replication> show
Properties:
             conflict_detected = true

Actions:

            TARGET          STATUS     NEXT
action-000  badenov         idle       Sync now
action-001  eel             idle       Sync now
action-002  bullwinkle      idle       Sync now

Children:
                         conflict => Conflict

goby:shares PRJ1 replication> conflict
goby:shares PRJ1 replication conflict> show
Properties:
                        source = (unset)

Sources:

SOURCE     HOSTNAME         ACTION_ID                              PROJECT
current    <current>
source-0   10.133.65.110    e99f0ea4-7328-488e-925e-8dccb3084b3d   PRJ1

goby:shares PRJ1 replication conflict> set source=source-0
                        source = source-0 (uncommitted)
goby:shares PRJ1 replication conflict> commit
goby:shares PJ1 replication conflict> 
```

Source: Resolving the Conflict using the BUI

1. On the incorrect source, select the project created by the unintended reverse.
2. Click **Replication**.

   The conflict notification should be displayed, as shown in the following figure.

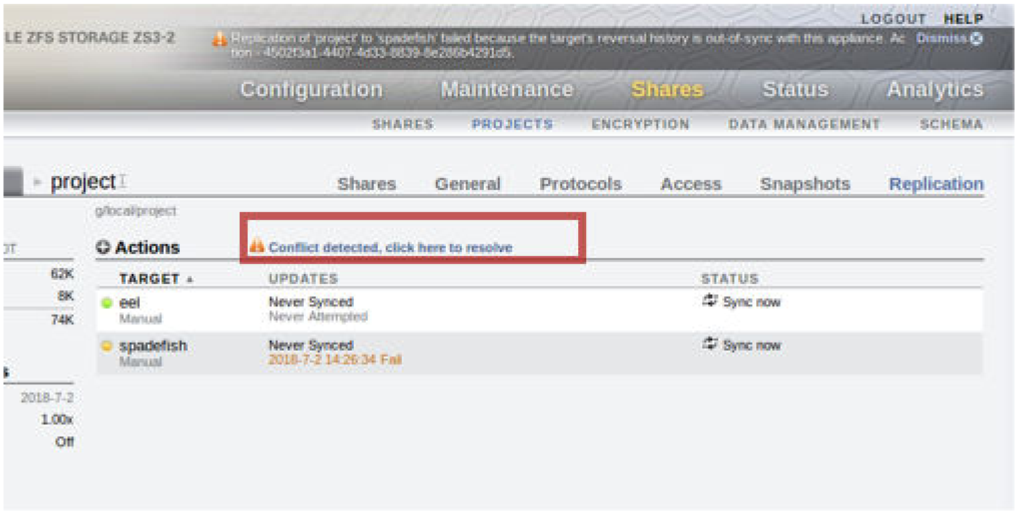

Replication of 'project' to 'spadefish' failed because the target's reversal history is out-of-sync with this appliance. Ac tion - 4502f3a1-4407-4d33-8839-8e286b4291d5.

Conflict detected, click here to resolve

TARGET	UPDATES	STATUS
eel Manual	Never Synced Never Attempted	Sync now
spadefish Manual	Never Synced 2018-7-2 14:26:34 Fail	Sync now

3. Click on the conflict notification to open the replication conflict resolution dialog box, as shown in the following figure.

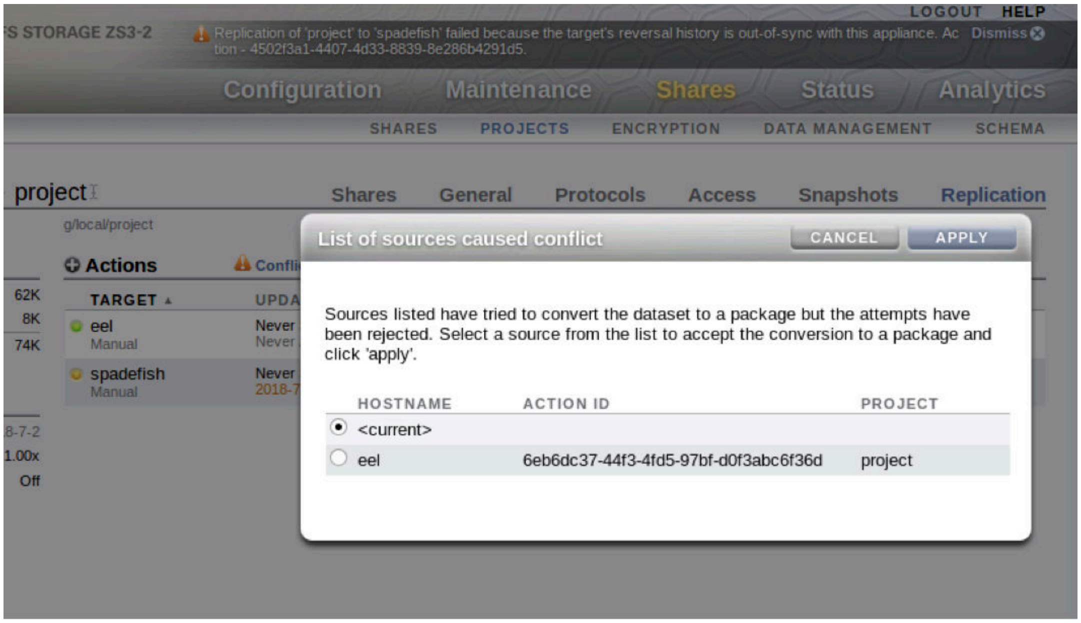

Replication of 'project' to 'spadefish' failed because the target's reversal history is out-of-sync with this appliance. Ac tion - 4502f3a1-4407-4d33-8839-8e286b4291d5.

List of sources caused conflict

Sources listed have tried to convert the dataset to a package but the attempts have been rejected. Select a source from the list to accept the conversion to a package and click 'apply'.

HOSTNAME	ACTION ID	PROJECT
● <current>		
○ eel	6eb6dc37-44f3-4fd5-97bf-d0f3abc6f36d	project

4. Select the correct source, and click **APPLY**.

   The very first update from that source will convert the project into a replication package, and the update from that source should succeed.

Compatibility Considerations

When configuring multi-target reversal, follow these compatibility rules:

- Only targets that support multi-target reversal or are running OS8.7.0 and later firmware can be members of a multi-target reversal group.
- Only targets that support multi-target reversal can be configured as potential sources.

- Only targets running OS8.7.0 and later firmware can be configured as dedicated targets in a multi-target reversal group.

The following failures occur due to incompatibility. This list is not exhaustive.

- An error occurs when configuring a potential source on a replication action to a target that does not support multi-target reversal.
- An error occurs when configuring a potential source under a dataset with existing replication actions to targets running firmware older than OS8.7.0.
- An error occurs when configuring a replication action to a target running firmware older than OS8.7.0 in an existing multi-target reversal group.
- Targets that are part of a multi-target reversal group and are configured as dedicated targets, but are running firmware older than OS8.8.6, can perform a replication reverse. However, they will not be able to perform a replication update back to their source. It is not possible to recover replication relations in such a scenario.

Automatic Snapshot Retention Implications

When automatic snapshot retention is used for some targets but not others in a multi-target reversal setup, after a reverse, targets that did not specify an automatic snapshot retention value keep the same number of snapshots as the new source, which might be different from what the targets had before. To avoid this situation, configure the retention policy for each target in a multi-target group as "independent." For more information, see Configuring Automatic Snapshot Retention on a Target (BUI) or Configuring Automatic Snapshot Retention on a Target (CLI).

Related Topics

- Disaster Recovery with Remote Replication
- Managing Replication Packages
- Reverse the Direction of Replication

Cascaded Replication

Cascaded replication allows replication of replicated data by configuring replication actions on replication packages. Multiple replication actions can be created on the same package in order to replicate the package to multiple target appliances.

Basic Setup

A basic cascaded replication setup consist of a source, one or more intermediate targets, and a final target.

The following example of cascaded replication consists of source node `A` that has an action set up on one of its projects to replicate it to a package on node `B`. Node `B` in turn has an action created to replicate this package to another package on node `C`, and node `C` has an action set up on this package to replicate it to a package on node `D`. Thus, node `B` and node `C` cascade a replica from node `A` to node `D`, as shown in the following figure.

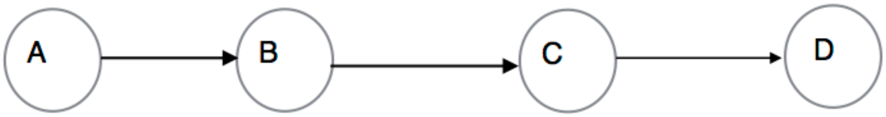

As shown in the following figure, source node A has an action set up on one of its projects to replicate it to a package on node B. Node B in turn has actions created to replicate this package to node C and node D. Thus, node B cascades a replica from node A to nodes C and D.

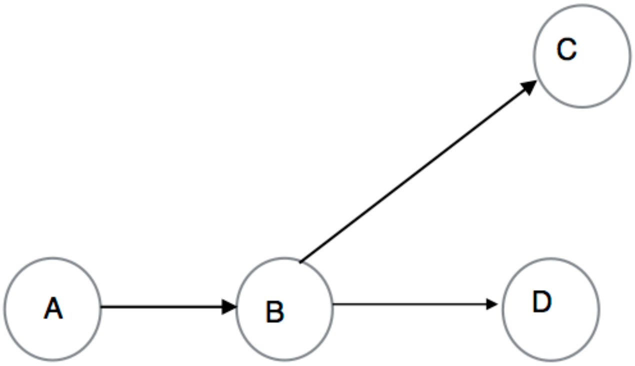

Replication reverse in a cascaded configuration is only supported on the nodes receiving replication updates directly from the source (node B in the previous example) and the reverse is disabled on any other nodes (nodes C and D). Multi-target reversal should be considered when planning for disaster recovery in a cascaded replication setup. See Multi-target Reversal and "Using the distant_target Property for Managing Reversal" in the following sections.

> **Note:**
>
> Replication reversal is not supported to a filesystem with mandatory file retention. For information about the file retention feature, see File Retention Management.

Cascaded Replication Management

Managing cascaded replication consists of two tasks:

- Creating an action under the replication package
- Configuring cascaded replication schedules

Creating an Action under the Replication Package

Similar to creating a replication action under a project, a replication action can be created under a replication package. The following figure shows how to create an action under a replication package using the BUI.

Note that cascaded replication is not supported for share-level replication. Actions cannot be configured under shares in replication packages, and actions cannot be configured under packages created as result of share-level replication.

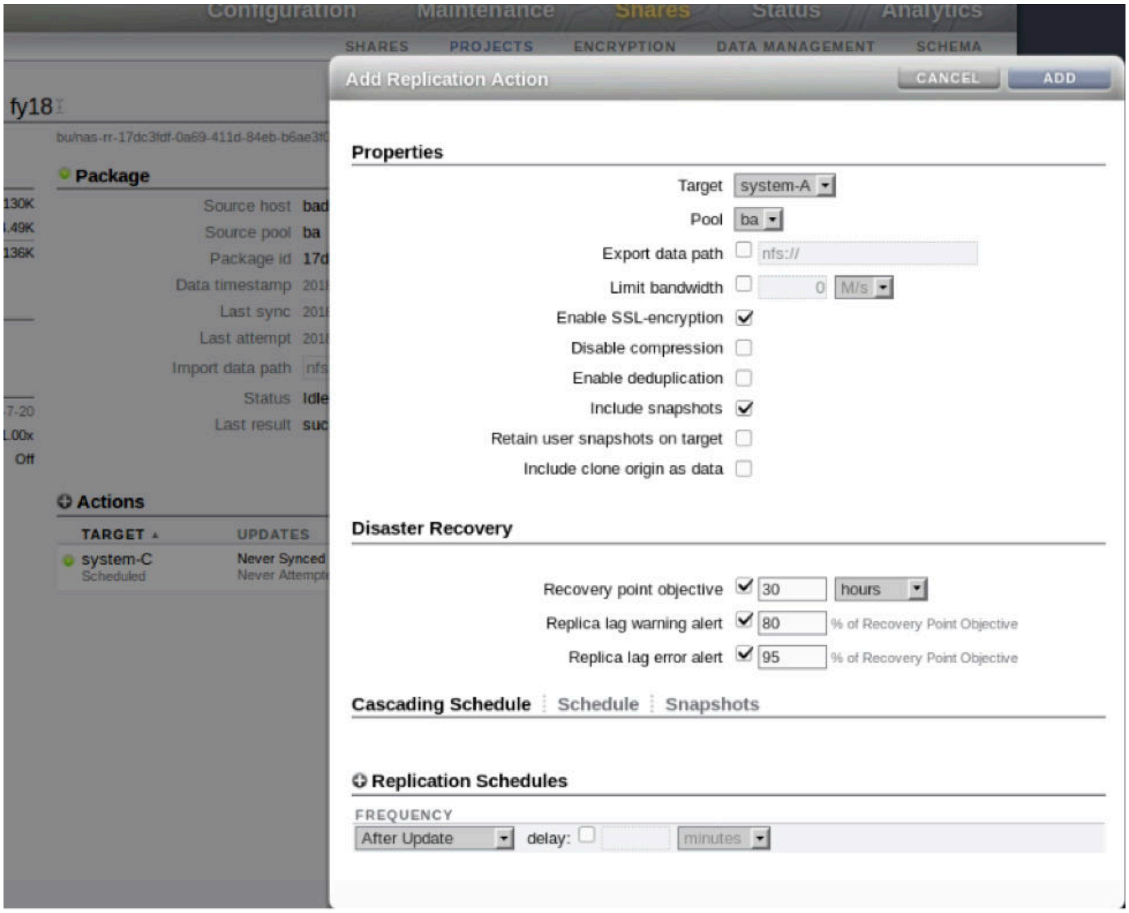

Add Replication Action

fy18

bu/nas-rr-17dc3fdf-0a69-411d-84eb-b6ae3f

Package

Source host bad
Source pool ba
Package id 17d
Data timestamp 201
Last sync 201
Last attempt 201
Import data path nfs
Status Idle
Last result suc

Actions

TARGET · UPDATES
system-C Scheduled · Never Synced Never Attempt

Properties

Target: system-A
Pool: ba
Export data path: □ nfs://
Limit bandwidth: □ 0 M/s
Enable SSL-encryption: ☑
Disable compression: □
Enable deduplication: □
Include snapshots: ☑
Retain user snapshots on target: □
Include clone origin as data: □

Disaster Recovery

Recovery point objective: ☑ 30 hours
Replica lag warning alert: ☑ 80 % of Recovery Point Objective
Replica lag error alert: ☑ 95 % of Recovery Point Objective

Cascading Schedule | Schedule | Snapshots

Replication Schedules

FREQUENCY
After Update delay: □ minutes

The following figure shows how to create an action under a replication package using the CLI.

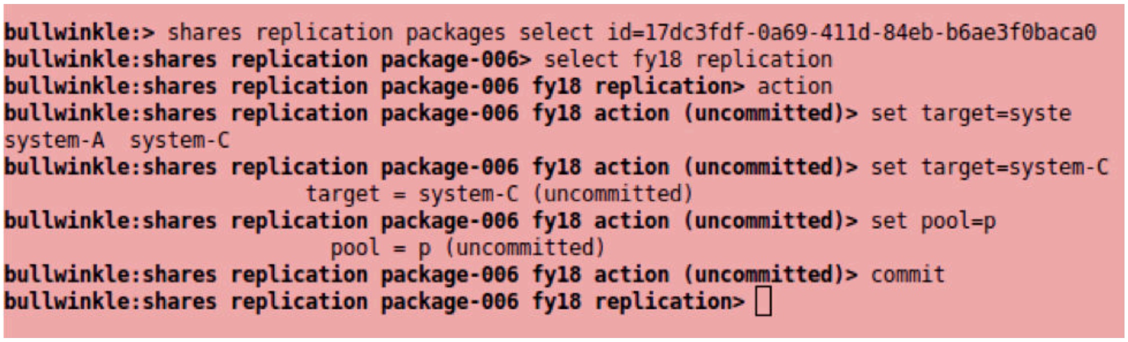

```
bullwinkle:> shares replication packages select id=17dc3fdf-0a69-411d-84eb-b6ae3f0baca0
bullwinkle:shares replication package-006> select fy18 replication
bullwinkle:shares replication package-006 fy18 replication> action
bullwinkle:shares replication package-006 fy18 action (uncommitted)> set target=syste
system-A  system-C
bullwinkle:shares replication package-006 fy18 action (uncommitted)> set target=system-C
                          target = system-C (uncommitted)
bullwinkle:shares replication package-006 fy18 action (uncommitted)> set pool=p
                            pool = p (uncommitted)
bullwinkle:shares replication package-006 fy18 action (uncommitted)> commit
bullwinkle:shares replication package-006 fy18 replication> 
```

Configuring Cascaded Replication Schedules

Every action has two schedules:

- A schedule on an action that is configured as part of a project is called a source schedule.

- A schedule on an action that is configured as part of a package is called a cascading schedule.

The following figure shows how to configure a new cascading schedule on a replication action using the CLI.

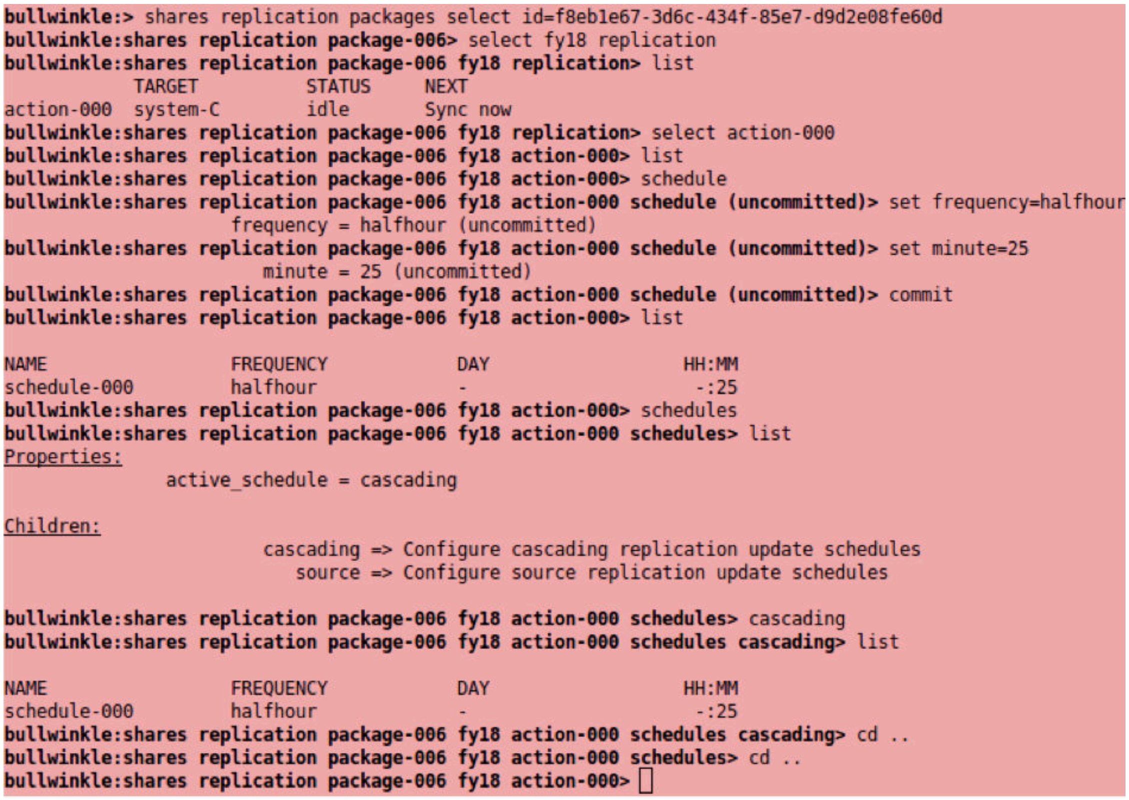

```
bullwinkle:> shares replication packages select id=f8eb1e67-3d6c-434f-85e7-d9d2e08fe60d
bullwinkle:shares replication package-006> select fy18 replication
bullwinkle:shares replication package-006 fy18 replication> list
            TARGET          STATUS     NEXT
action-000  system-C        idle       Sync now
bullwinkle:shares replication package-006 fy18 replication> select action-000
bullwinkle:shares replication package-006 fy18 action-000> list
bullwinkle:shares replication package-006 fy18 action-000> schedule
bullwinkle:shares replication package-006 fy18 action-000 schedule (uncommitted)> set frequency=halfhour
                     frequency = halfhour (uncommitted)
bullwinkle:shares replication package-006 fy18 action-000 schedule (uncommitted)> set minute=25
                        minute = 25 (uncommitted)
bullwinkle:shares replication package-006 fy18 action-000 schedule (uncommitted)> commit
bullwinkle:shares replication package-006 fy18 action-000> list

NAME                 FREQUENCY            DAY                  HH:MM
schedule-000         halfhour             -                     -:25
bullwinkle:shares replication package-006 fy18 action-000> schedules
bullwinkle:shares replication package-006 fy18 action-000 schedules> list
Properties:
               active_schedule = cascading

Children:
                        cascading => Configure cascading replication update schedules
                           source => Configure source replication update schedules

bullwinkle:shares replication package-006 fy18 action-000 schedules> cascading
bullwinkle:shares replication package-006 fy18 action-000 schedules cascading> list

NAME                 FREQUENCY            DAY                  HH:MM
schedule-000         halfhour             -                     -:25
bullwinkle:shares replication package-006 fy18 action-000 schedules cascading> cd ..
bullwinkle:shares replication package-006 fy18 action-000 schedules> cd ..
bullwinkle:shares replication package-006 fy18 action-000> 
```

The following figure shows how to configure a new cascading schedule on a replication action using the BUI.

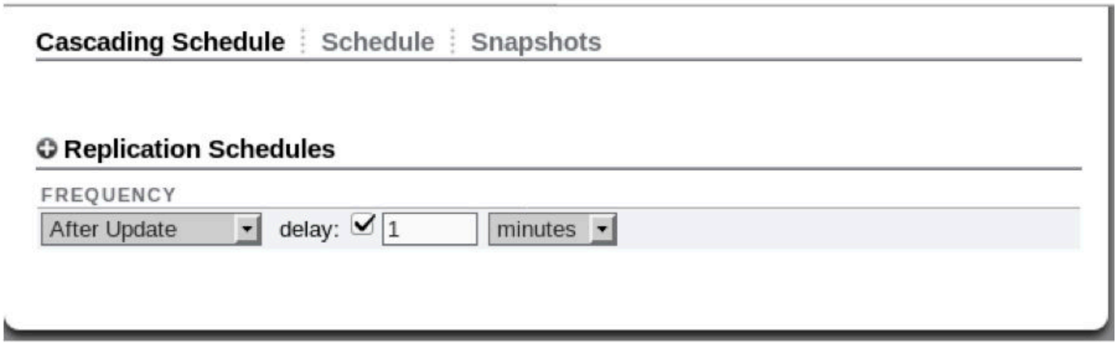

Cascading Schedule | Schedule | Snapshots

Replication Schedules

FREQUENCY

After Update ▾ delay: ☑ 1 minutes ▾

Both schedules can be configured at any time, but only one schedule is effective, depending where the action is configured. If there is a change in replication topology due to replication reverse or conversion, the effective schedule may change. For example, after performing a reverse on a replication package, actions configured in the reversed package are now configured in the project, and so their source schedules become active. Similarly, after performing a replication conversion on project actions that used to be configured under the project are now configured in the package, and so their cascading schedules are active.

The following figure shows how to view the effective schedule using the CLI.

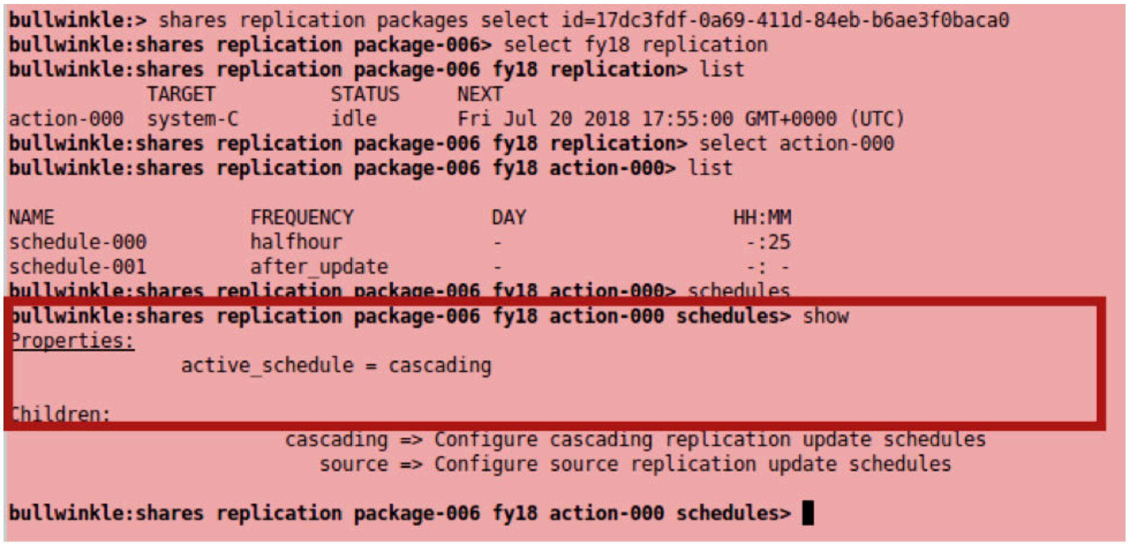

```
bullwinkle:> shares replication packages select id=17dc3fdf-0a69-411d-84eb-b6ae3f0baca0
bullwinkle:shares replication package-006> select fy18 replication
bullwinkle:shares replication package-006 fy18 replication> list
            TARGET          STATUS     NEXT
action-000  system-C        idle       Fri Jul 20 2018 17:55:00 GMT+0000 (UTC)
bullwinkle:shares replication package-006 fy18 replication> select action-000
bullwinkle:shares replication package-006 fy18 action-000> list

NAME                 FREQUENCY            DAY                  HH:MM
schedule-000         halfhour             -                     -:25
schedule-001         after_update         -                     -: -
bullwinkle:shares replication package-006 fy18 action-000> schedules
bullwinkle:shares replication package-006 fy18 action-000 schedules> show
Properties:
               active_schedule = cascading

Children:
                        cascading => Configure cascading replication update schedules
                           source => Configure source replication update schedules

bullwinkle:shares replication package-006 fy18 action-000 schedules> █
```

The following figure shows how to view the effective schedule using the BUI. The leftmost schedule tab on an action is the effective schedule.

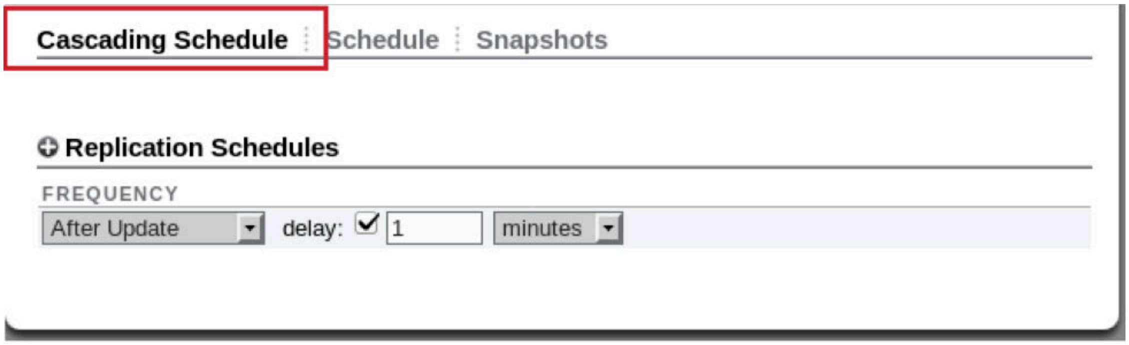

FREQUENCY

After Update ▾ delay: ☑ 1 minutes ▾

Both source and cascading schedules are preserved when new actions are created during replication reverse. See Reverse the Direction of Replication and Multi-target Reversal.

The `after_update` option is available for cascading schedules. When set, it will initiate the update only after an incoming update has completed.

The `update_cascade_delay` property can be used to specify the delay between the completion of an incoming replication update and the initiation of the outgoing cascading update. This property is effective only when the `after_update` option is selected.

The `after_update` option and the `update_cascade_delay` property can be configured as shown in the following figure.

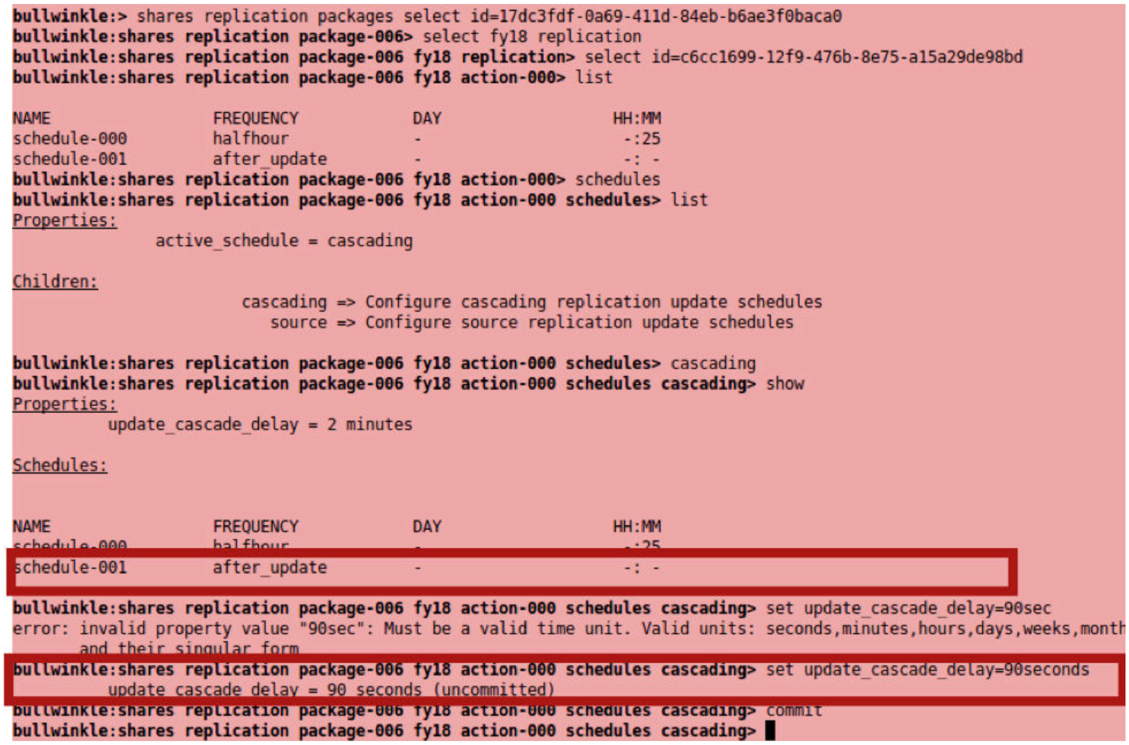

```
bullwinkle:> shares replication packages select id=17dc3fdf-0a69-411d-84eb-b6ae3f0baca0
bullwinkle:shares replication package-006> select fy18 replication
bullwinkle:shares replication package-006 fy18 replication> select id=c6cc1699-12f9-476b-8e75-a15a29de98bd
bullwinkle:shares replication package-006 fy18 action-000> list

NAME                 FREQUENCY            DAY                  HH:MM
schedule-000         halfhour             -                     -:25
schedule-001         after_update         -                     -: -
bullwinkle:shares replication package-006 fy18 action-000> schedules
bullwinkle:shares replication package-006 fy18 action-000 schedules> list
Properties:
               active_schedule = cascading

Children:
                        cascading => Configure cascading replication update schedules
                           source => Configure source replication update schedules

bullwinkle:shares replication package-006 fy18 action-000 schedules> cascading
bullwinkle:shares replication package-006 fy18 action-000 schedules cascading> show
Properties:
          update_cascade_delay = 2 minutes

Schedules:


NAME                 FREQUENCY            DAY                  HH:MM
schedule-000         halfhour             -                     -:25
schedule-001         after_update         -                     -: -

bullwinkle:shares replication package-006 fy18 action-000 schedules cascading> set update_cascade_delay=90sec
error: invalid property value "90sec": Must be a valid time unit. Valid units: seconds,minutes,hours,days,weeks,month
       and their singular form
bullwinkle:shares replication package-006 fy18 action-000 schedules cascading> set update_cascade_delay=90seconds
          update_cascade_delay = 90 seconds (uncommitted)
bullwinkle:shares replication package-006 fy18 action-000 schedules cascading> commit
bullwinkle:shares replication package-006 fy18 action-000 schedules cascading> █
```

Note that the continuous schedule option is not available for actions under packages, and cascading replication actions do not create new replication snapshots. When a new update is delivered from the source, cascading actions use the latest package snapshot to perform replication updates to its targets. Therefore, performing an update on a cascading action without first receiving an update from the source will have no effect; the successful update status will be set and the `last_try` property will be updated.

Using the distant_target Property for Managing Reversal

In cascaded replication, a distant target relationship between two nodes differentiates between local and distant targets. The `distant_target` property preserves this relationship after topology transformations as result of performing replication reverse or conversion.

The following configuration scenarios illustrate the `distant_target` property:

1. When the `distant_target` property is not set and there is no multi-target reversal set up. In this case, when a reverse is performed on a target, not all actions will be preserved. For more information, see Multi-target Reversal. As shown in the following figure, neither Nodes `A` nor `B` have a distant target relationship, nor Node `B` is configured as a potential source (action `AB`: `potential_source=false`, `distant_target=false`).

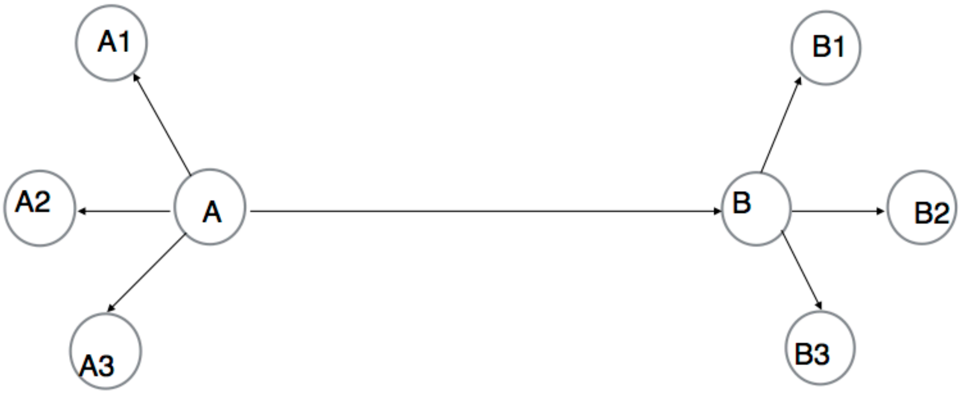

When A fails and a reverse on B is performed, no actions to A{1..3} are created, as shown in the following figure.

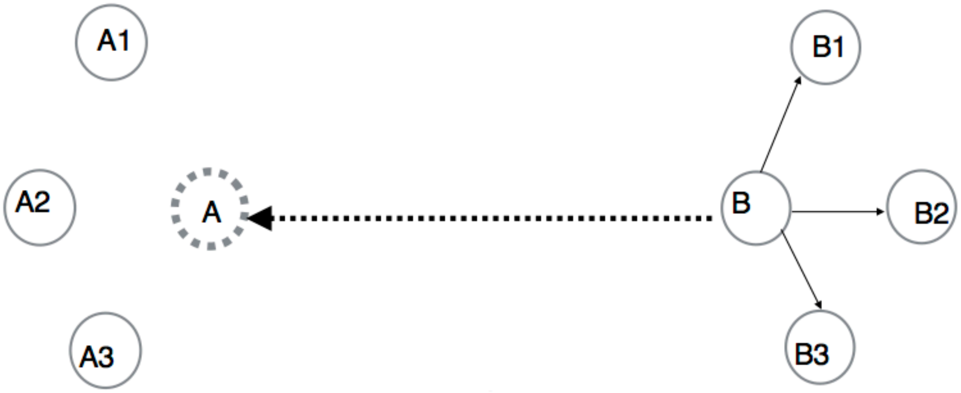

2. When the distant_target property is not set and there is multi-target reversal set up. In this case, when a reverse is performed on a potential source, it creates actions to all targets in the multi-target reversal group. There is no distant target relationship and, therefore, there is no differentiation between local and distant targets. Node A is the source and Node B is the potential source, as shown in the following figure. Nodes A and B have no distant target relationship (action AB: potential_source=true, distant_target=false).

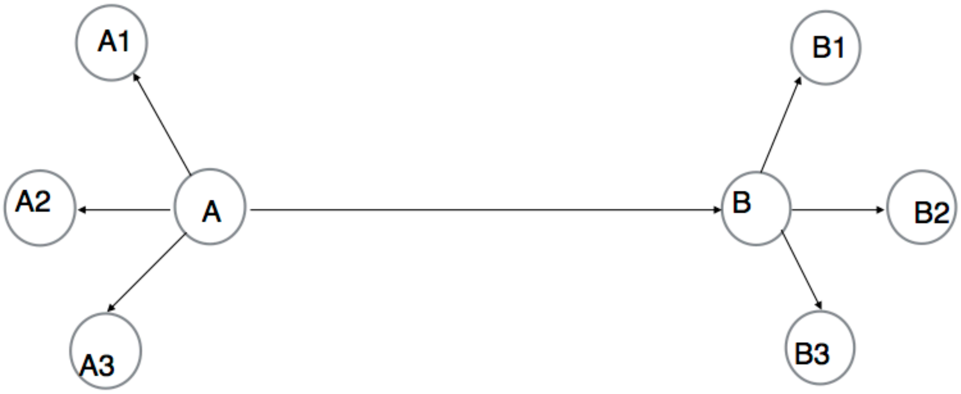

However, when Node A fails and Node B is the new source after a reverse is performed, actions to A{1..3} are created, as shown in the following figure.

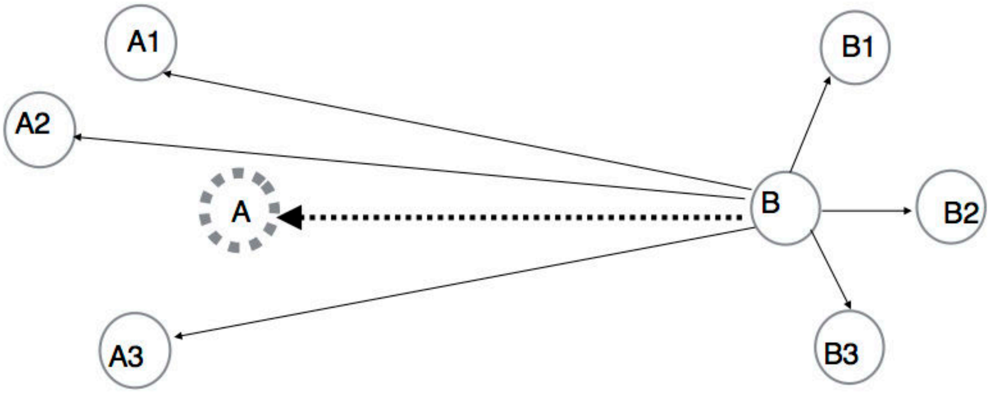

3. When the distant_target property is set and there is a multi-target reversal set up. In this case, when a reverse is performed on a potential source, it does not create actions to the targets that are local to the distant source. Node A is the source and Node B is the potential source, as shown in the following figure. Nodes A and B have a distant target relationship (action AB: potential_source=true, distant_target=true).

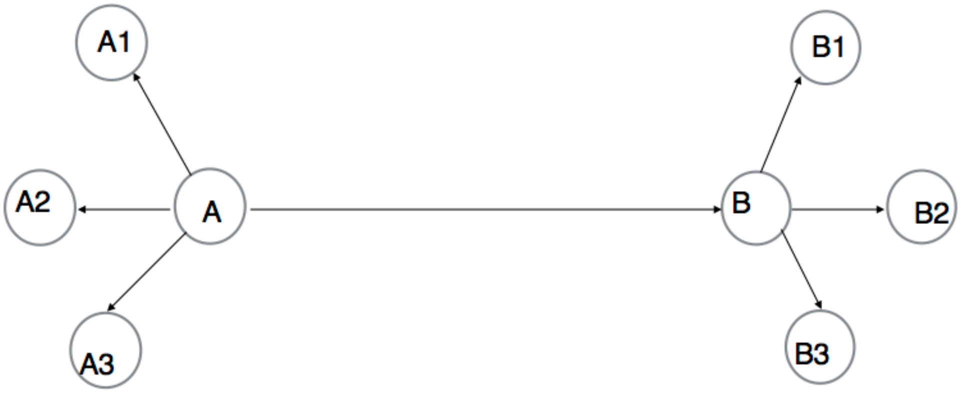

When Node `A` fails and after a reverse on Node `B`, Node `B` becomes the new source. However, because of the distant target relationship between `A` and `B`, actions to `A{1..3}` are not created. After `A` recovers, it will cascade its package to `A{1..3}`.

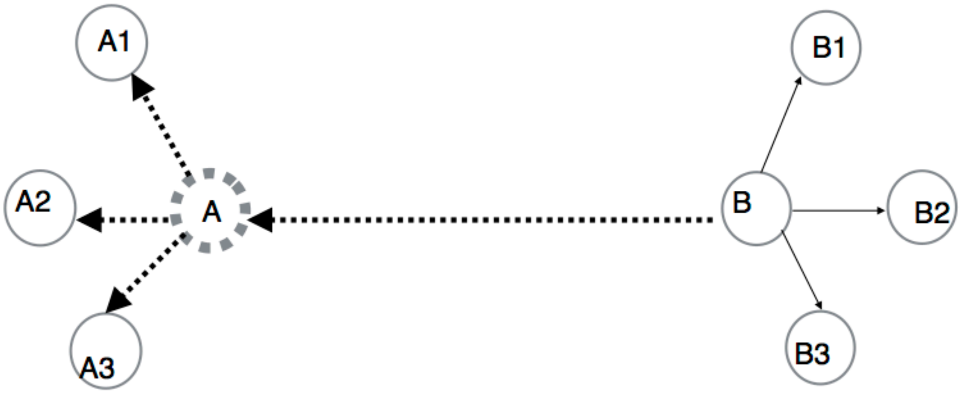

The following figure shows how to configure a distant target using the CLI.

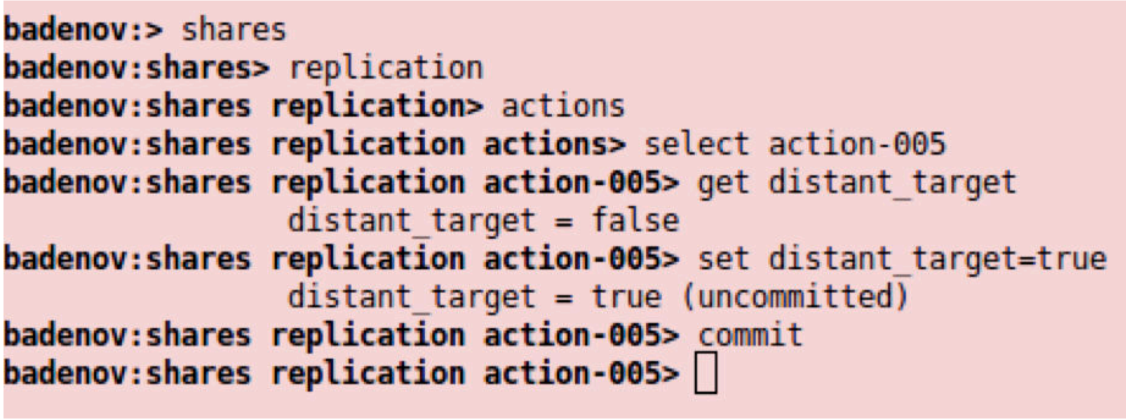

```
badenov:> shares
badenov:shares> replication
badenov:shares replication> actions
badenov:shares replication actions> select action-005
badenov:shares replication action-005> get distant_target
                   distant_target = false
badenov:shares replication action-005> set distant_target=true
                   distant_target = true (uncommitted)
badenov:shares replication action-005> commit
badenov:shares replication action-005> 
```

The following figure shows how to configure a distant target using the BUI.

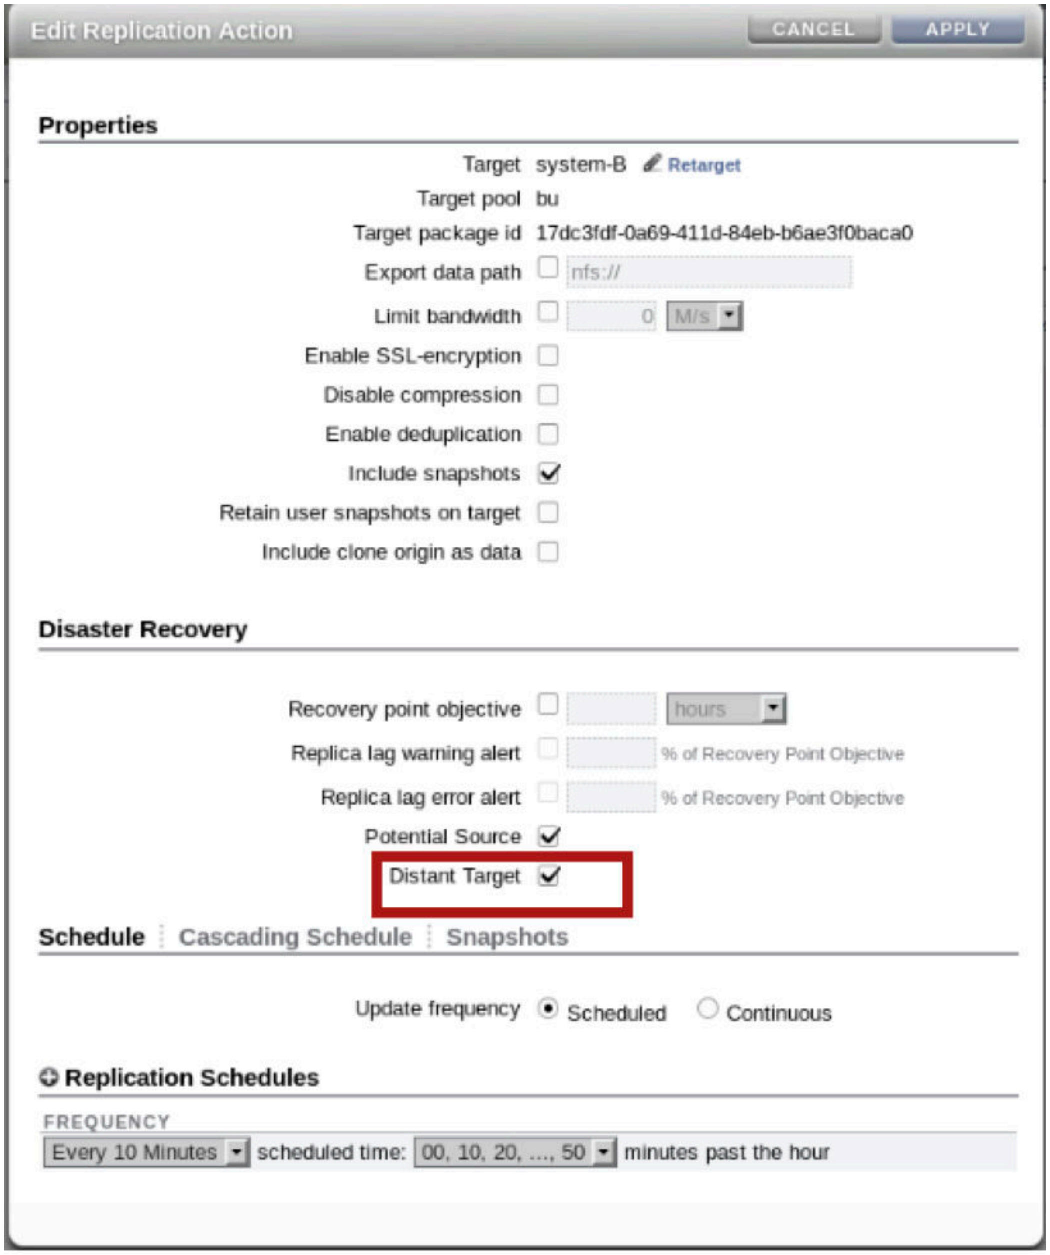

Properties

Target system-B Retarget
Target pool bu
Target package id 17dc3fdf-0a69-411d-84eb-b6ae3f0baca0
Export data path □ nfs://
Limit bandwidth □ 0 M/s
Enable SSL-encryption □
Disable compression □
Enable deduplication □
Include snapshots ☑
Retain user snapshots on target □
Include clone origin as data □

Disaster Recovery

Recovery point objective □ hours
Replica lag warning alert □ % of Recovery Point Objective
Replica lag error alert □ % of Recovery Point Objective
Potential Source ☑
Distant Target ☑

Update frequency ● Scheduled ○ Continuous

Replication Schedules

FREQUENCY
Every 10 Minutes scheduled time: 00, 10, 20, ..., 50 minutes past the hour

It is recommend to configure multi-target reversal by setting a potential source when using cascaded replication; otherwise, it may not be possible to return to the original cascaded configuration after multiple replication reversals. Also, the distant target can only be set when a potential source is set. For more information, see Multi-target Reversal.

Failure Scenarios

Failure of a single node in a cascaded replication can impact multiple appliances or even multiple data centers connected to each other by replication relations. Although there is no limitation on the cascaded replication topology, the following typical failure scenarios are presented:

- Original source failure
- Intermediate target failure
- Source site failure

Original Source Failure

In this scenario, the source of the cascading chain fails. Recovery from this failure requires an initial multi-target reversal configuration on the source. The process is initiated on a target by performing a replication reverse. The reversed target becomes the source and continues to send incremental updates to all other original targets and to the original source. For further details, see Multi-target Reversal.

As shown in the following figure, `A` is the source and `A1` is the potential source. `A` is replicated to `A1` and `B`. `B`, in turn, replicates to `C` and `D`, and source `A` fails.

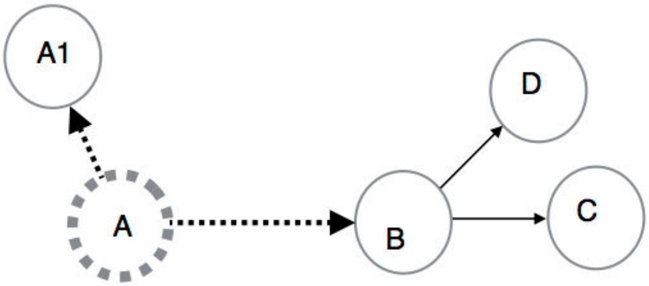

As shown in the following figure, after a reverse is performed on `A1`, it replicates to `A` (when `A` recovers) and to `B` which, in turn, continues to send updates to `C` and `D`.

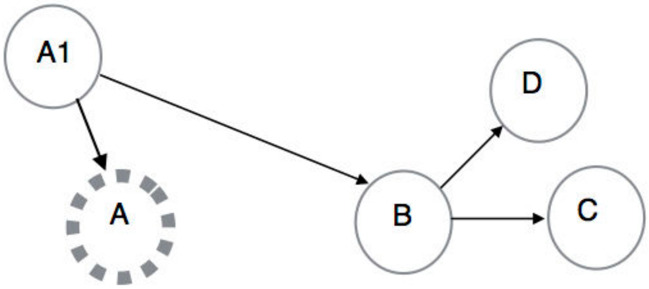

Intermediate Target Failure

This failure scenario pertains to the failure of an intermediate node in a cascaded chain. Recovery from this scenario can be achieved by following the retargeting procedure. As shown in the following figure, intermediate node `B` in the cascaded chain fails.

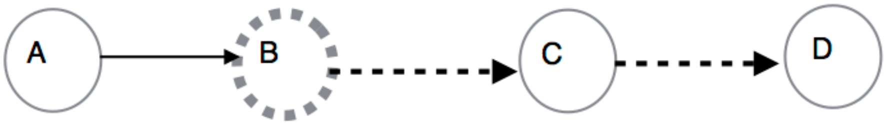

As shown in the following figure, after the retarget/bypass procedure is applied, source `A` bypasses failed node `B` and continues sending updates to node `C`, which, in turn, sends an update to node `D`.

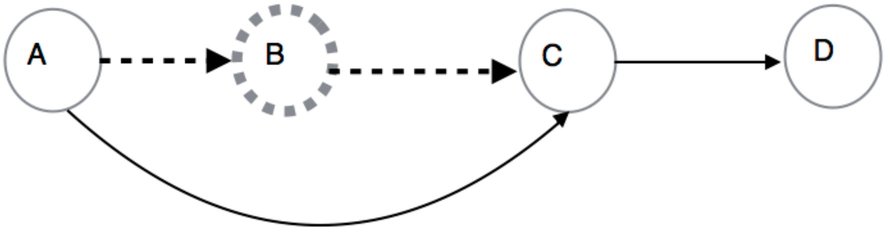

As shown in the following figure, when B recovers, action between A and B is restored, while action between B and C remains bypassed.

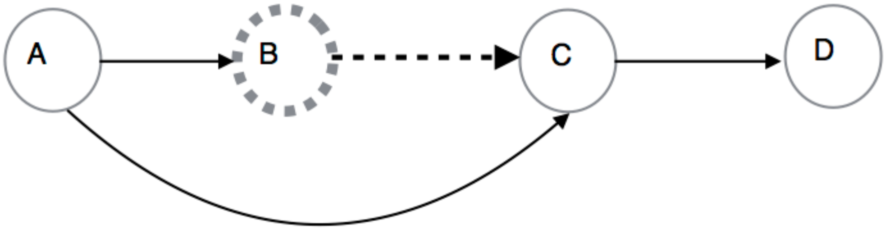

Bypassing a failed node using the CLI:

1. Select the action to the failed node on the source.
2. Enter `retarget`.
3. Set the `retarget_mode` property to `bypass`.
4. Set the target to be the bypass target, as shown in the following figure.

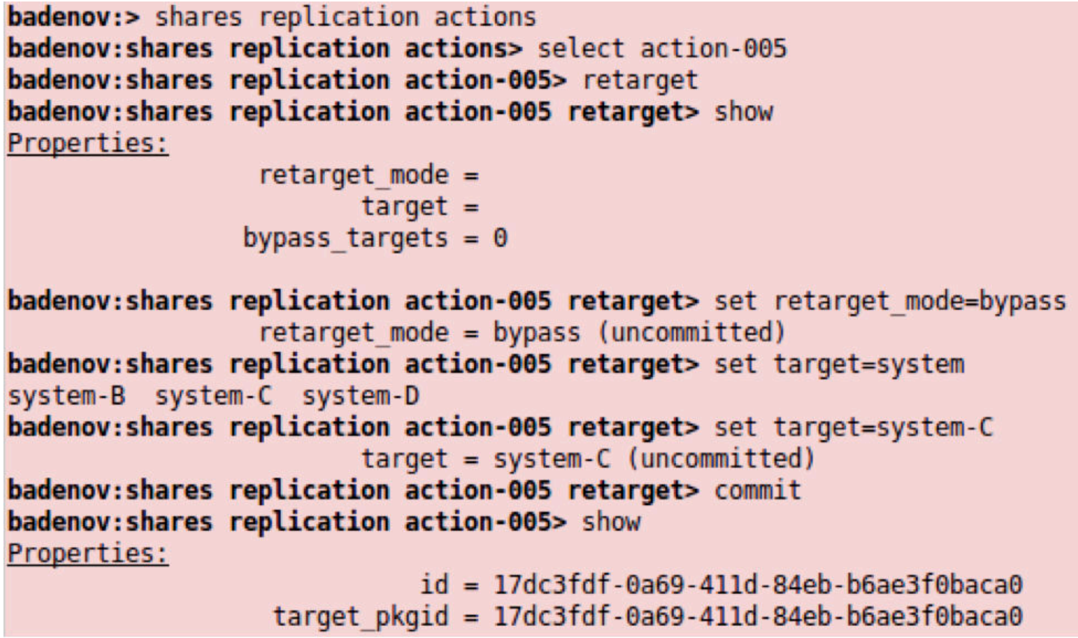

```
badenov:> shares replication actions
badenov:shares replication actions> select action-005
badenov:shares replication action-005> retarget
badenov:shares replication action-005 retarget> show
Properties:
                 retarget_mode =
                        target =
                bypass_targets = 0

badenov:shares replication action-005 retarget> set retarget_mode=bypass
                 retarget_mode = bypass (uncommitted)
badenov:shares replication action-005 retarget> set target=system
system-B  system-C  system-D
badenov:shares replication action-005 retarget> set target=system-C
                        target = system-C (uncommitted)
badenov:shares replication action-005 retarget> commit
badenov:shares replication action-005> show
Properties:
                            id = 17dc3fdf-0a69-411d-84eb-b6ae3f0baca0
                  target_pkgid = 17dc3fdf-0a69-411d-84eb-b6ae3f0baca0
```

5. The `bypassed_id` property shown in the bypassing action can be used to select the original action to perform a retarget restore after the failed node has recovered. The `bypassed_id` generated can be viewed, as shown in the following figure.

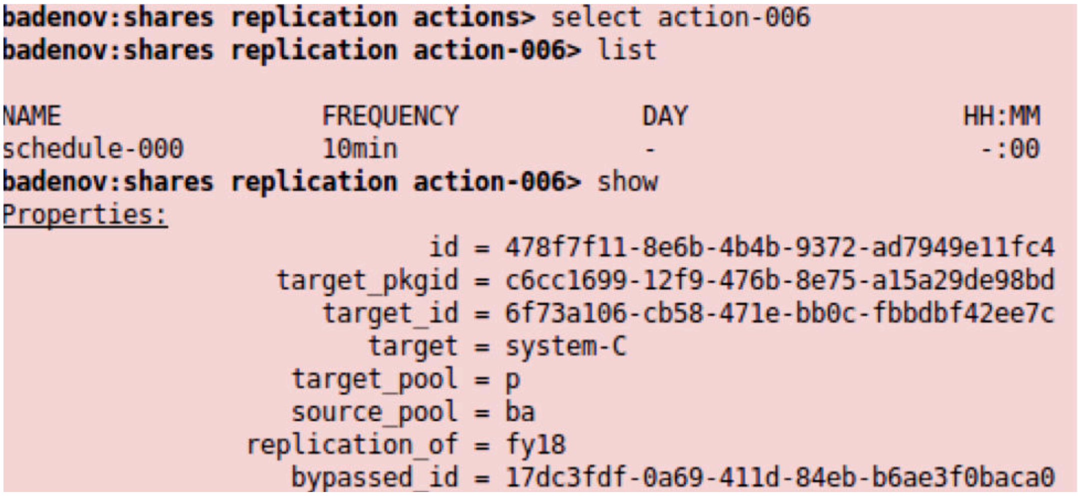

```
badenov:shares replication actions> select action-006
badenov:shares replication action-006> list

NAME                    FREQUENCY             DAY                   HH:MM
schedule-000            10min                 -                     -:00
badenov:shares replication action-006> show
Properties:
                              id = 478f7f11-8e6b-4b4b-9372-ad7949e11fc4
                    target_pkgid = c6cc1699-12f9-476b-8e75-a15a29de98bd
                       target_id = 6f73a106-cb58-471e-bb0c-fbbdbf42ee7c
                          target = system-C
                     target_pool = p
                     source_pool = ba
                  replication_of = fy18
                     bypassed_id = 17dc3fdf-0a69-411d-84eb-b6ae3f0baca0
```

6. At a later stage when `B` has recovered, use the retarget/restore procedure to get back to the original cascaded topology, as shown in the following figure.

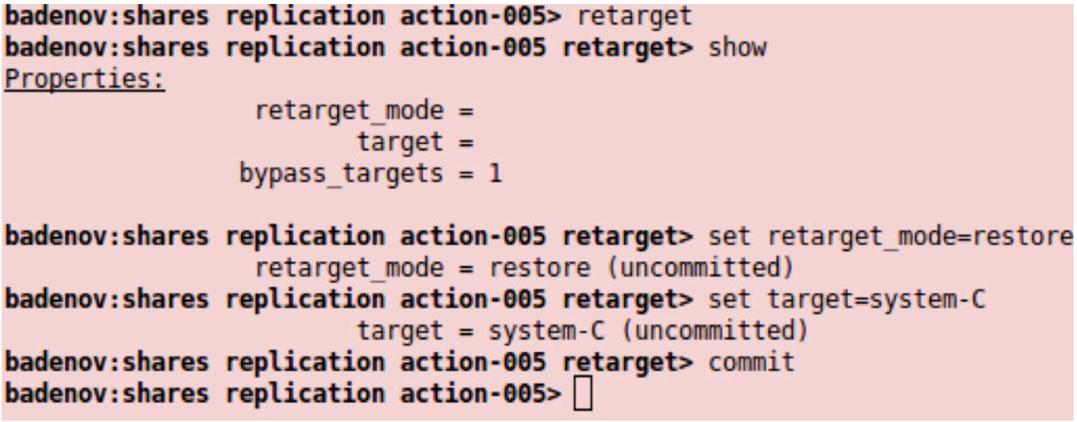

```
badenov:shares replication action-005> retarget
badenov:shares replication action-005 retarget> show
Properties:
                 retarget_mode =
                        target =
                bypass_targets = 1

badenov:shares replication action-005 retarget> set retarget_mode=restore
                 retarget_mode = restore (uncommitted)
badenov:shares replication action-005 retarget> set target=system-C
                        target = system-C (uncommitted)
badenov:shares replication action-005 retarget> commit
badenov:shares replication action-005>
```

Bypassing a failed node using the BUI:

1. On the source node, select the action to the failed node.
2. In the **Edit Replication Action** dialog box, click **Retarget**.
3. In the **Replication Retarget** dialog box, set the mode to **Bypass**.
4. Select the target from the drop-down list to be the bypass target, as shown in the following figure.

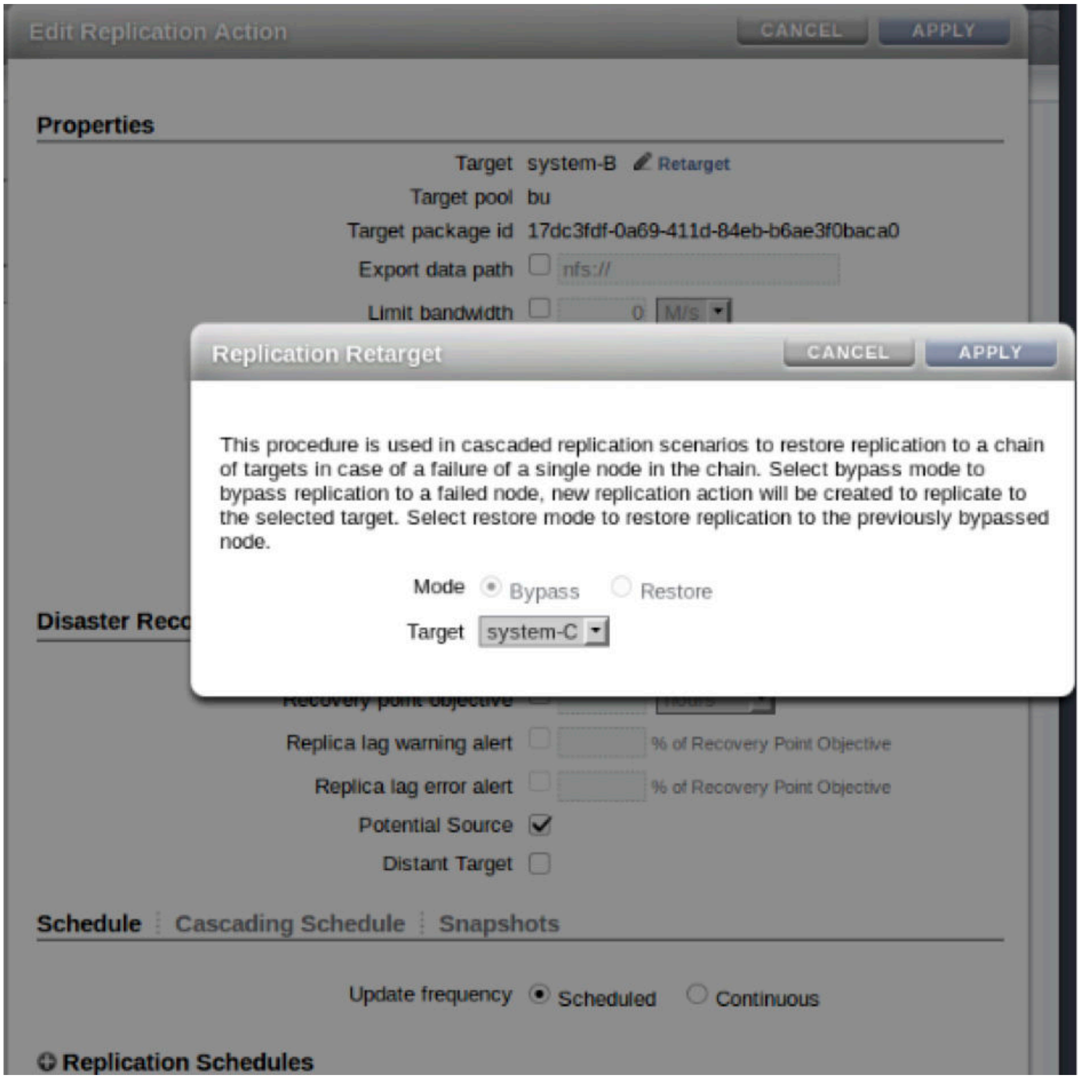

Target system-B Retarget

Target pool bu

Target package id 17dc3fdf-0a69-411d-84eb-b6ae3f0baca0

Export data path □ nfs://

Limit bandwidth □ 0 M/s

This procedure is used in cascaded replication scenarios to restore replication to a chain of targets in case of a failure of a single node in the chain. Select bypass mode to bypass replication to a failed node, new replication action will be created to replicate to the selected target. Select restore mode to restore replication to the previously bypassed node.

Mode ● Bypass ○ Restore

Target system-C

Recovery point objective □ hours

Replica lag warning alert □ % of Recovery Point Objective

Replica lag error alert □ % of Recovery Point Objective

Potential Source ☑

Distant Target □

Update frequency ● Scheduled ○ Continuous

5. Click **APPLY**.
6. At a later stage when B has recovered, use the restore procedure to get back to the original cascaded topology, as shown in the following figure.

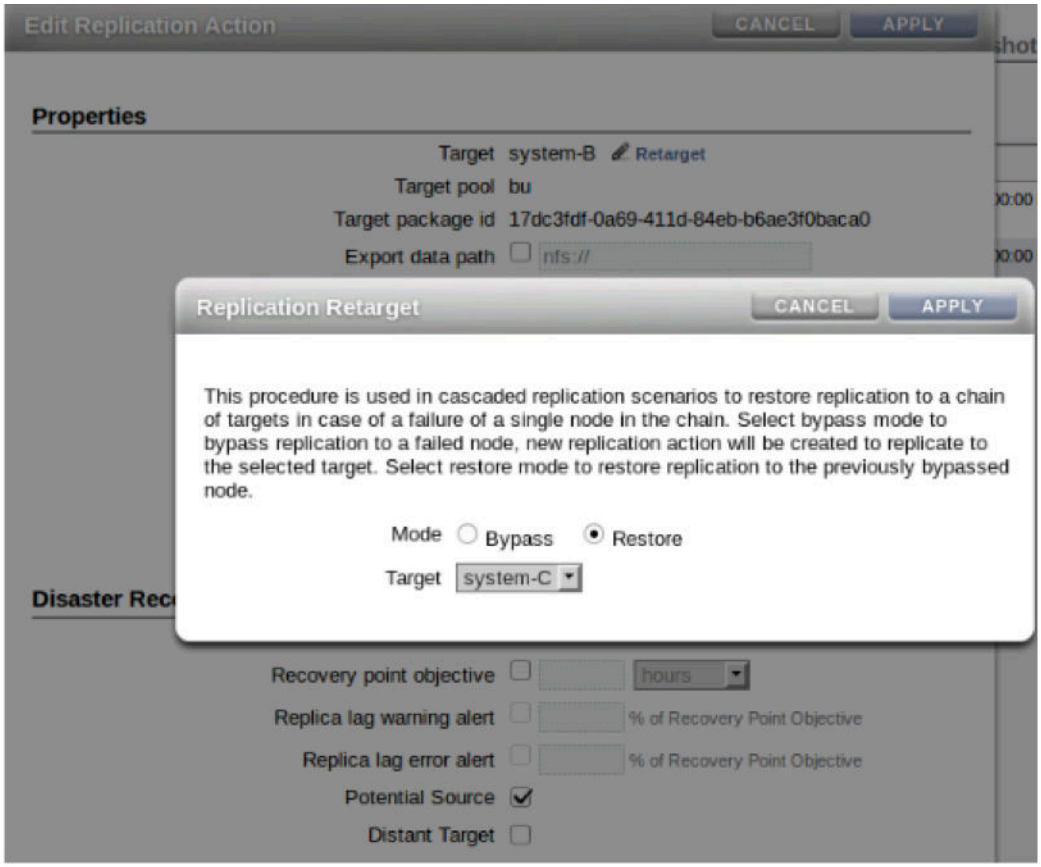

Properties

Target system-B Retarget

Target pool bu

Target package id 17dc3fdf-0a69-411d-84eb-b6ae3f0baca0

Export data path □ nfs://

This procedure is used in cascaded replication scenarios to restore replication to a chain of targets in case of a failure of a single node in the chain. Select bypass mode to bypass replication to a failed node, new replication action will be created to replicate to the selected target. Select restore mode to restore replication to the previously bypassed node.

Mode ○ Bypass ● Restore

Target system-C

Disaster Rec

Recovery point objective □ hours

Replica lag warning alert □ % of Recovery Point Objective

Replica lag error alert □ % of Recovery Point Objective

Potential Source ☑

Distant Target □

Source Site Failure

In this scenario, the source and all its direct local targets are not available, and are replaced by a new source at one of the distant targets. Recovery from this scenario requires multi-target reversal configuration on the source. For further details, see Multi-target Reversal. It also requires some targets to be configured as distant targets. For details, refer to "Using the distant_target Property for Managing Reversal."

As shown in the following figure, `A0`, `B0` and `C0` form a multi-target reversal group, where `B0` and `C0` are potential sources and distant targets. `A0,B0` and `A0,C0` form a distant target relation (`A0B0` and `A0C0` should be configured as potential source and distant target), while `A1,A2,A3` form a local relation with `A0`, and `B1,B2,B3` form a local relation with `B0`.

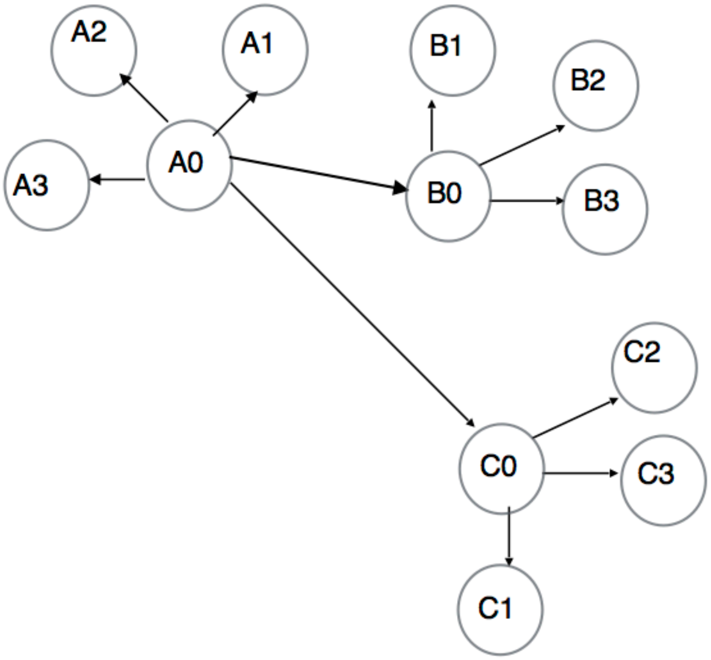

In the case when a reverse is performed on `B0`, as shown in the following figure, it treats `A1,A2,A3` as local targets of `A0` and hence does not recreate actions to `A1,A2,A3`. Due to the properties of multi-target reversal, `B0` sends updates to `C0`.

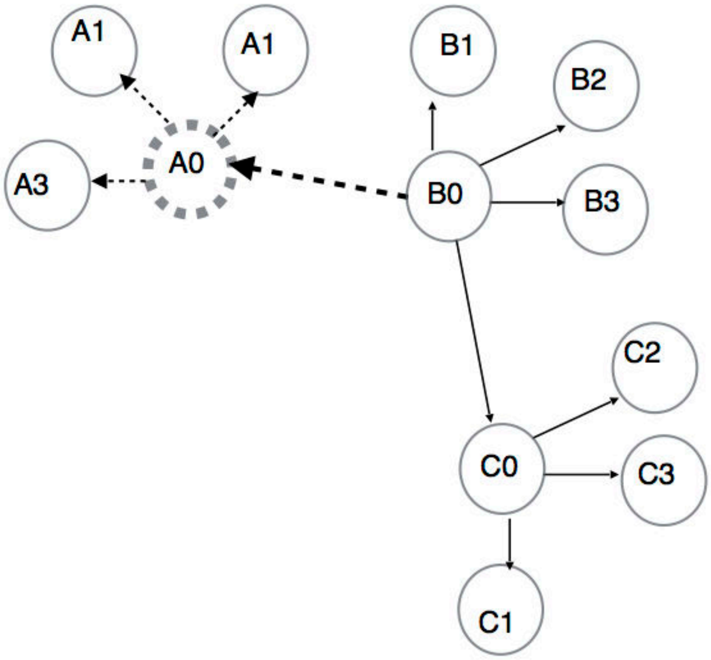

Compatibility Considerations

When configuring cascaded replication, follow these compatibility rules:

- Nodes running firmware older than OS8.7.0 cannot be configured as part of cascaded replication.
- Nodes running firmware later than OS8.7.0 can be configured as final nodes.
- Final nodes running firmware that does not support cascaded replication may perform replication reverse. However, they will not be able to perform a replication update back to their source. It is not possible to recover replication relations in such a scenario.
- The retarget cannot be performed to final nodes running firmware that does not support cascaded replication.
- Replication compatibility rules for targets running different firmware in cascaded chains are the same as for any source and target. For example, when a target does not support a feature used by the source, the replication update from the corresponding source may fail.
- System rollback to a firmware version that does not support cascaded replication will result in the destruction of all cascaded replication actions.

Related Topics

- Replication Action Properties
- Multi-target Reversal

Exporting Replicated Filesystems

Replicated filesystems can be exported read-only to NAS clients. This can be used to verify the replicated data or to perform backups or other intensive operations on the replicated data (offloading such work from the source appliance).

The filesystem's contents always match the most recently received replication snapshot for that filesystem. This may be newer than the most recently received snapshot for the entire package, and it may not match the most recent snapshot for other shares in the same package. For details, see Replication Snapshots and Data Consistency.

Replication updates are applied atomically at the filesystem level. Clients looking at replicated files will see replication updates as an instantaneous change in the underlying filesystem. Clients working with files deleted in the most recent update will see errors. Clients working with files changed in the most recent update will immediately see the updated contents.

Replicated filesystems are not exported by default. They are exported by modifying the exported property of the project or share using the BUI or CLI, as shown in the following figure.

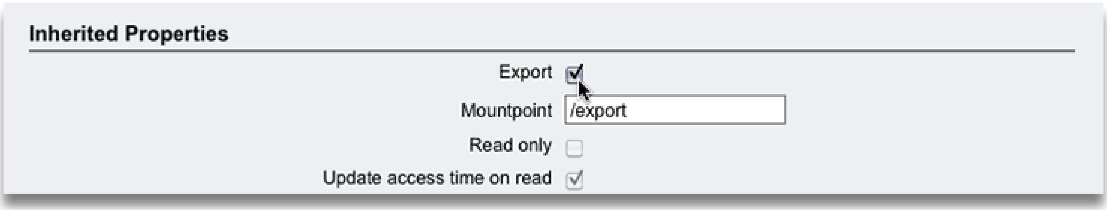

Inherited Properties

Export ☑
Mountpoint /export
Read only ☐
Update access time on read ☑

This property is inherited like other share properties. This property is not shown for local projects and shares because they are always exported. Additionally, severing replication (which converts the package into a local project) causes the package's shares to become exported.

Replicated LUNs currently cannot be exported. They must be first cloned or the replication package severed in order to export their contents.

Related Topics

- Remote Replication Workflow
- Inherited Properties

Severing Replication

A replication package can be converted into a local, writable project that behaves just like other local projects (that is, without the management restrictions applied to replication packages) by severing the replication connection. Severing a replication package can be used to migrate data between appliances or in other scenarios that do not involve replicating the received data back to the source appliance.

If a replication update is performed during or after a sever operation, the update will fail with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target.

A new replication action and a full update of the same project is required to send replication updates to a new replication package.

To avoid mount point or SMB name conflicts, resolve the conflicts before severing the replication package by reconfiguring the project, share mount points, or SMB resource names. Because all local shares are always exported, and might be shared over SMB, the sever operation will fail if any mount points or SMB resource names conflict between replicated filesystems and other filesystems on the system.

Related Topics

- Severing a Replication Package - BUI, CLI
- Disaster Recovery with Remote Replication
- Managing User-Generated Snapshots

Reverse the Direction of Replication

The direction of the replication can be reversed to support two-system disaster recovery plans and disk-to-disk backups.

Reversing Replication for Disaster Recovery

The reverse replication operation converts the replication package into a local project. This operation also configures a replication action on the new local project for incremental replication back to the source appliance. The first update attempt will convert the original project on the source system into a replication package and roll back any changes made since the last successful replication update from that system.

The following figure describes a typical reverse replication sequence of events.

Using Remote Replication for Disaster Recovery

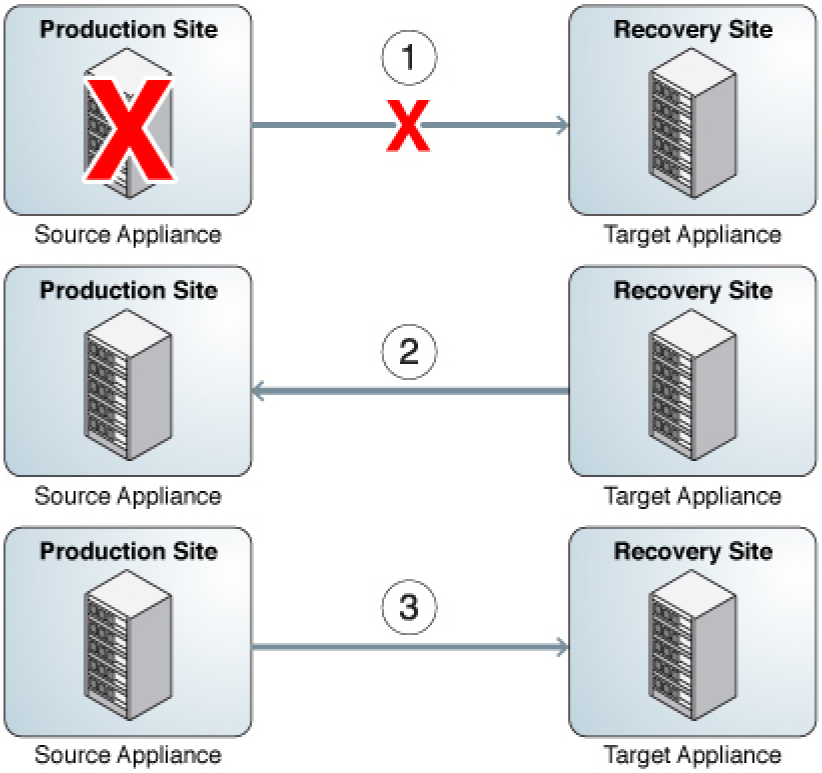

1. The production system is the source appliance serving the client workload and replicating to the replication target located at a recovery site.

   A complete failure of the source appliance occurs at the production site. From the recovery site, the administrator reverses the direction of replication. This operation converts the replication package to a local, writable project. The administrator redirects client workloads and failover IP addresses to the recovery site.

2. After the production site is restored and back to normal operations, the administrator initiates a replication update from the recovery site to the production site.

   This operation converts the production copy into a replication package, and rolls back any changes written to the recovery site while the production site was down.

3. After the production site is updated, the administrator reverses the direction of replication again, which makes the copy at the production site writable.

   The administrator then redirects client workloads and failover IP address back to the production site. The original relationship between the source appliance at the production site and the replication target at the recovery site is restored.

Share-level and Project-level Reversal

When the original source project is converted into a replication package on the original source appliance (which is now acting as the target), the shares that were replicated as part of the action/package currently being reversed are moved into a new replication package and unexported. The original project remains in the local collection, but may end up empty if the action/package included all of its shares. When share-level replication is reversed, any other shares in the original project remain unchanged.

Before reversing the direction of replication for a package, stop replication updates of that project from the source appliance. If a replication update is in progress when an administrator reverses the direction of replication for a project, administrators cannot know which consistent replication snapshot was used to create the resulting project on the former replication target (now source appliance).

If a replication update is performed during or after a reversal operation, the update fails with an appropriate alert. The replication action is then disabled, resulting in no future updates from this action to the replication target. A new replication action and a full update are required to send updates from the original project to a new replication package.

Because all local shares are exported, all shares in a package are exported when the package is reversed, whether or not they were previously exported. If there are mount point conflicts between replicated filesystems and other filesystems on the system, the reverse operation will fail. These conflicts must be resolved before severing by reconfiguring the mount points of the relevant shares. Because this operation is typically part of the critical path of restoring production service, it is strongly recommended to resolve these mount point conflicts when the systems are first set up rather than at the time of disaster recovery failover.

> **Note:**
>
> Replication reversal is not supported to a filesystem with mandatory file retention nor to its project. For information about the file retention feature, see File Retention Management.

Related Topics

- Disaster Recovery with Remote Replication
- Managing Replication Packages
- Multi-target Reversal

Destroying a Replication Package

The project and shares within a package cannot be destroyed without destroying the entire package. The entire package can be destroyed from the BUI by destroying the corresponding project. A package can be destroyed from the CLI using the `destroy` command at the `shares replication packages` node.

When a package is destroyed, subsequent replication updates from the corresponding action will fail. To resume replication, the action will need to be recreated on the source to create a new package on the target into which to receive a new copy of the data.

Target Replica Backups

You can back up target replica datasets (projects or shares) using the NDMP zfs backup type. Replica backup is enabled on the appliance by applying the deferred update "Support for NDMP zfs-type Replica Backup." The replica backup feature chooses the most recent system-generated snapshot to be backed up, unless you specify a user-generated (non .rr extension) snapshot. For more information, see Replica Backups.

Some older replication snapshots, those originally preserved for future incremental backups, might not be needed and can be deleted. If the snapshot is held by NDMP, a confirmation is displayed warning of the potential impact to ongoing or future NDMP backups.

The following sequence of events causes a replication failure and generates an alert. For information about recovering from this error, see "Name collision" in Replication Failures.

1. A replica snapshot is held by NDMP on the replication target (for an ongoing backup or a future incremental backup).
2. The corresponding snapshot on the source appliance is deleted or renamed.
3. A new snapshot is created on the source appliance with the same name as the replica snapshot held on the replication target.
4. A replication update is attempted.

Raw Crypto Replication

Raw crypto replication improves the security and efficiency of replicating encrypted data by avoiding both decrypting the data on the source appliance and reencrypting it on the target appliance. Security is enhanced by sending the data encrypted instead of unencrypted.

The raw crypto option is enabled by default for new replication actions if the raw crypto replication deferred update has been applied. Once enabled, raw crypto replication can be disabled before the initial update.

The Raw Crypto Replication Deferred Update is available in Oracle ZFS Storage software version OS8.8.57 or later. To use raw crypto replication, both the source and target appliances must have accepted the deferred update. It is recommended to accept the deferred update on the targets first to avoid action update delays. Updates will fail if the source appliance's replication action uses the raw crypto feature but the deferred update has not been accepted on the target appliance. For information on applying deferred updates, see Deferred Updates in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

Additional restrictions: 1) Both the source and target appliances must have identical encryption keys in each dataset's `keyname` property (same name and same contents); and 2) When using raw crypto replication with the multi-target reversal feature, all members must use raw crypto replication.

To disable the raw crypto property for a new replication action, select check box **Disable rawcrypto mode** in the BUI or set `rawcrypto` to `off` in CLI node `shares` *`project`* `action-`*`number`*.

Related Topics

- Creating a Replication Action (BUI)
- Creating a Replication Action (CLI)
- Replication Action Properties

8
Data Encryption

> **Note:**
>
> Encryption is a licensed feature for certain models. For details, refer to the "Oracle Software License Agreement ("SLA") and Entitlement for Hardware Systems with Integrated Software Options" and the Licensing Information User Manual for the software release.

Oracle ZFS Storage Appliance offers transparent data encryption for pools and for individual shares (filesystems and LUNs) and shares created inside of projects.

To configure and manage encryption, use these tasks:

- Data Encryption Workflow
- Configuring LOCAL Keystore Encryption - BUI, CLI
- Configuring OKM Keystore Encryption - BUI, CLI
- Configuring KMIP Keystore Encryption - BUI, CLI
- Creating an Encrypted Pool - BUI, CLI
- Changing a Pool Encryption Key - BUI, CLI
- Creating an Encrypted Project - BUI, CLI
- Changing a Project Encryption Key - BUI, CLI
- Creating an Encrypted Filesystem or LUN - BUI, CLI
- Changing a Share Encryption Key - BUI, CLI
- Backing up a LOCAL Key - BUI, CLI
- Deleting an Encryption Key - BUI, CLI
- Restoring a LOCAL Key - BUI, CLI
- Cloning a Snapshot - BUI, CLI

To understand data encryption, use these topics:

- Encryption Properties
- Managing Encryption Keys
- Performance Impact of Encryption
- Encryption Key Life Cycle
- Backing Up, Replicating, and Restoring Encrypted Projects and Shares

Data Encryption Workflow

The following steps show the general procedure for configuring and using data encryption. For information about encryption properties, see Encryption Properties.

1. Configure LOCAL keystore, Oracle Key Manager (OKM) keystore, or Key Management Interoperability Protocol (KMIP) keystore encryption.
   - Configuring LOCAL Keystore Encryption - BUI, CLI
   - Configuring OKM Keystore Encryption - BUI, CLI
   - Configuring KMIP Keystore Encryption - BUI, CLI
2. Create encryption keys and certificates.

   Use the same procedures as in the preceding step.
3. Optional: Create a pool using an encryption key.

   For information, see Creating an Encrypted Pool - BUI, CLI.
4. Optional: Create a project using an encryption key.

   For information, see Creating an Encrypted Project - BUI, CLI.
5. Create a share in a project that uses an encryption key or create a share using an encryption key.

   For information about creating a share, see Shares and Projects or Creating an Encrypted Project - BUI, CLI.

Related Topics

- Managing Encryption Keys
- Performance Impact of Encryption
- Encryption Key Life Cycle
- Backing Up, Replicating, and Restoring Encrypted Projects and Shares

Configuring LOCAL Keystore Encryption (BUI)

To configure encryption using the LOCAL keystore, first set up the master passphrase, and then create keys. For information about encryption properties, see Encryption Properties.

1. From the **Shares** menu, select **Encryption**.
2. Click **Local**.

   The LOCAL keystore information is displayed.

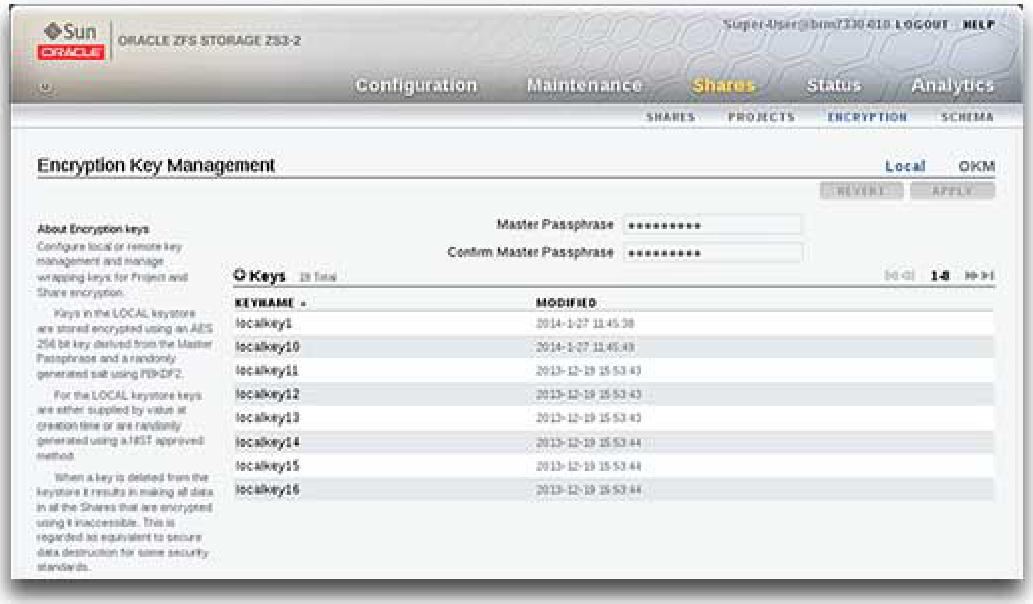

Encryption Key Management

About Encryption keys

Configure local or remote key management and manage wrapping keys for Project and Share encryption.

Keys in the LOCAL keystore are stored encrypted using an AES 256 bit key derived from the Master Passphrase and a randomly generated salt using PBKDF2.

For the LOCAL keystore keys are either supplied by value at creation time or are randomly generated using a NIST approved method.

When a key is deleted from the keystore it results in making all data in all the Shares that are encrypted using it inaccessible. This is regarded as equivalent to secure data destruction for some security standards.

Master Passphrase ••••••••••

Confirm Master Passphrase ••••••••••

Keys 19 Total

KEYNAME	MODIFIED
localkey1	2014-1-27 11:45:38
localkey10	2014-1-27 11:45:49
localkey11	2013-12-19 15:53:43
localkey12	2013-12-19 15:53:43
localkey13	2013-12-19 15:53:43
localkey14	2013-12-19 15:53:44
localkey15	2013-12-19 15:53:44
localkey16	2013-12-19 15:53:44

3. To configure the master passphrase, type the passphrase supplied by your administrator, and then retype it in the next box.
4. To save the master passphrase, click **APPLY**.
5. To create a key, click the add icon ⊕ next to **Keys**.

   The **New Key** dialog box is displayed.

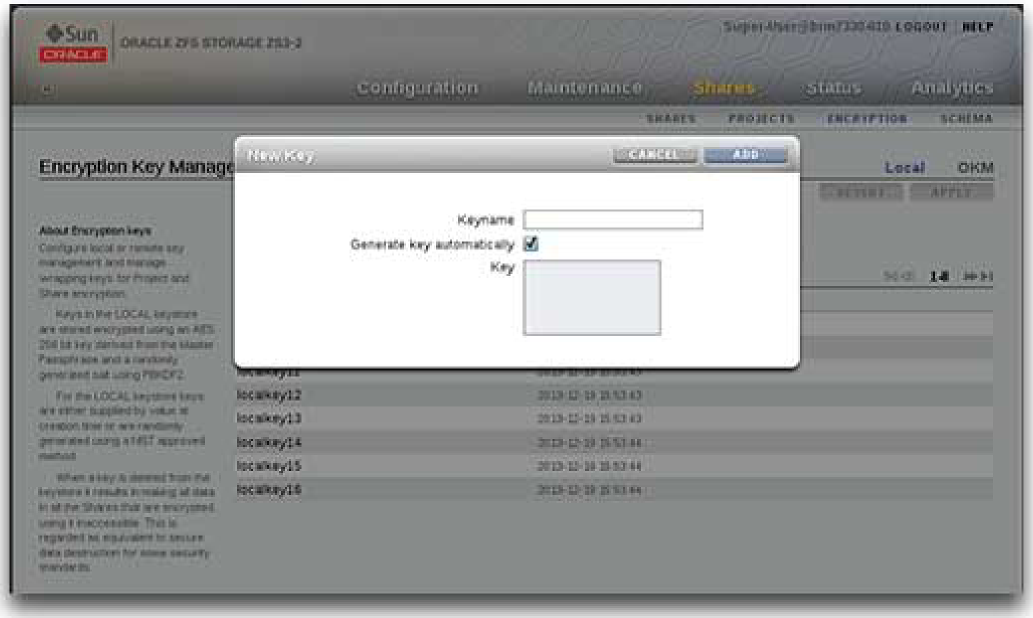

New Key

Keyname

Generate key automatically ☑

Key

6. Enter the key name.
7. Optional: Enter a key.

A key is generated automatically if you leave the **Generate key automatically** box checked.

To provide your own key, do the following:

a. Uncheck the **Generate key automatically** box.

b. Enter your hex-encoded, raw 256-bit key in the **Key** box.

8. To save the key, click **ADD**.

   When you click **ADD**, the new key appears in the list of keys with the creation date.

Related Topics

- Configuring OKM Keystore Encryption (BUI)
- Configuring KMIP Keystore Encryption (BUI)
- Configuring LOCAL Keystore Encryption (CLI)
- Creating an Encrypted Pool (BUI)
- Creating an Encrypted Project (BUI)

Configuring LOCAL Keystore Encryption (CLI)

To configure encryption using the LOCAL keystore, first set up the master passphrase, and then create keys. For information about encryption properties, see Encryption Properties.

1. Configure LOCAL keystore encryption.

   To configure LOCAL keystore encryption, set the master passphrase.

```
hostname:> shares encryption
hostname:shares encryption> show
Children:
                            local => Manage LOCAL keystore
                              okm => Register keys with Oracle Key Manager

hostname:shares encryption> local
hostname:shares encryption local> show
Properties:
             master_passphrase =

Children:
                       keys => Manage this Keystore's Keys

hostname:shares encryption local> set master_passphrase
Enter new master_passphrase:
Re-enter new master_passphrase:
             master_passphrase = (set) (uncommitted)
b7420-16m:shares encryption local> commit
b7420-16m:shares encryption local> show
Properties:
             master_passphrase = (set)

Children:
                       keys => Manage this Keystore's Keys
```

2. Create a LOCAL key.

   To create a key, enter a key name. The value of the `keyname` property is the name used in the CLI and BUI when assigning a key to a pool, project, or share.

You can either leave the `key` property empty, and the system will generate a random key value, or you can enter a hex-encoded, raw 256-bit key value. In the following example, the system generates the `key` value.

The key is stored in an encrypted form using the master passphrase from step 1.

```
hostname:shares encryption local> keys
hostname:shares encryption local keys> show
Keys:

NAME     CREATED                CIPHER KEYNAME
Properties:
                        cipher = AES
                           key =
                       keyname = (unset)
hostname:shares encryption local> create
hostname:shares encryption local key-000 (uncommitted)> set keyname=Key-0
                       keyname = Key-0 (uncommitted)
hostname:shares encryption local key-000 (uncommitted)> commit
hostname:shares encryption local keys> show
Keys:

NAME     CREATED                CIPHER KEYNAME
key-000  2019-7-1 18:43:33      AES    Key-0
hostname:shares encryption local keys> select key-000
hostname:shares encryption local key-000> show
Properties:
                        cipher = AES
                           key =
ce968122d0bba26c3d66b6985ee358d18a786607f80eb4ebd834e4404fe8aa84
                       keyname = Key-0
```

Related Topics

- Configuring OKM Keystore Encryption (CLI)
- Configuring KMIP Keystore Encryption (CLI)
- Configuring LOCAL Keystore Encryption (BUI)
- Creating an Encrypted Pool (CLI)
- Creating an Encrypted Project (CLI)

Configuring OKM Keystore Encryption (BUI)

To configure encryption using Oracle Key Manager (OKM), first set up the key manager server information, and then create keys. For information about encryption properties, see Encryption Properties.

> **Note:**
>
> If the Oracle ZFS Storage Appliance system is clustered, do not use the "one time passphrase" setting when creating the OKM server agent. If you use the "one time passphrase" setting in this situation, registration on the other cluster node will fail and keys will not be available on failover.

1. From the **Shares** menu, select **Encryption**.

2. Click **OKM**.

   The OKM keystore information is displayed.

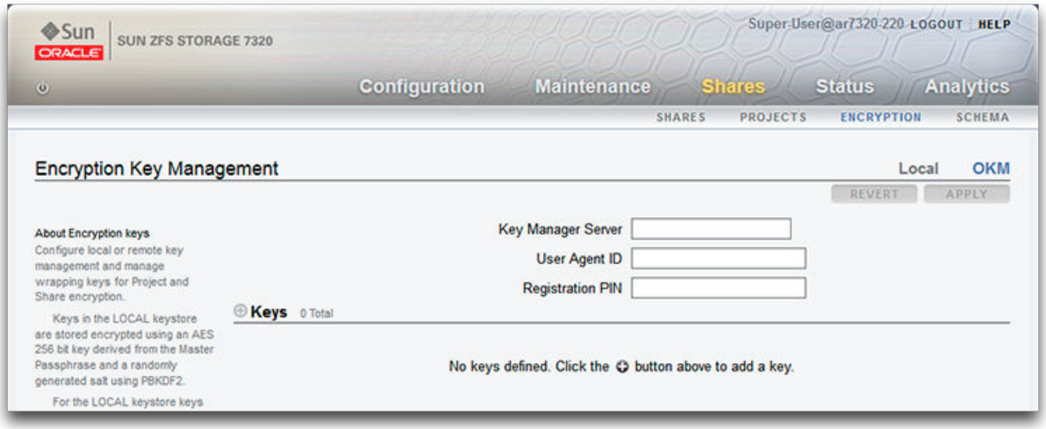

Encryption Key Management

Key Manager Server

User Agent ID

Registration PIN

About Encryption keys
Configure local or remote key management and manage wrapping keys for Project and Share encryption.

Keys in the LOCAL keystore are stored encrypted using an AES 256 bit key derived from the Master Passphrase and a randomly generated salt using PBKDF2.

For the LOCAL keystore keys

No keys defined. Click the ⊕ button above to add a key.

3. To configure the server information, type the following information:
   - **Key Manager Server**
   - **User Agent ID**
   - **Registration PIN**
4. To save the server information, click **APPLY**.
5. To create a key, click the add icon ⊕ next to **Keys**.

   The **New Key** dialog box is displayed.
6. Type a name for the key.
7. To save the key, click **ADD**.

   When you click **ADD**, the new key appears in the list of keys with the creation date.

Related Topics

- Oracle Key Manager (OKM) Keystore
- Configuring LOCAL Keystore Encryption (BUI)
- Configuring KMIP Keystore Encryption (BUI)
- Configuring OKM Keystore Encryption (CLI)
- Creating an Encrypted Pool (BUI)
- Creating an Encrypted Project (BUI)

Configuring OKM Keystore Encryption (CLI)

To configure encryption using OKM, first set up the key manager server information, and then create keys. For information about encryption properties, see Encryption Properties.

> **Note:**
>
> If the Oracle ZFS Storage Appliance system is clustered, do not use the "one time passphrase" setting when creating the OKM server agent. If you use the "one time passphrase" setting in this situation, registration on the other cluster node will fail and keys will not be available on failover.

1. Configure OKM keystore encryption.

   To configure OKM keystore encryption, set the agent ID, registration PIN (supplied by your OKM security officer), and server IP address.

   ```
   hostname:> shares encryption
   hostname:shares encryption> show
   Children:
                               local => Manage LOCAL keystore
                                 okm => Register keys with Oracle Key Manager

   hostname:shares encryption> okm
   hostname:shares encryption okm> show
   Properties:
                        agent_id = ExternalClient041
                registration_pin = (set)
                     server_addr = 10.80.180.109

   Children:
                               keys => Manage this Keystore's Keys
   ```

2. Create an OKM key.

   To create a key, set a key name.

   ```
   hostname:shares encryption okm keys> create
   hostname:shares encryption okm key-372 (uncommitted)> ls
   Properties:
                          cipher = AES
                         keyname = (unset)
   hostname:shares encryption okm key-372 (uncommitted)> set keyname=anykey
                         keyname = anykey (uncommitted)
   hostname:shares encryption okm key-372 (uncommitted)> commit
   ```

Related Topics

- Oracle Key Manager (OKM) Keystore
- Configuring LOCAL Keystore Encryption (CLI)
- Configuring KMIP Keystore Encryption (CLI)
- Configuring OKM Keystore Encryption (BUI)
- Creating an Encrypted Pool (CLI)
- Creating an Encrypted Project (CLI)

Configuring KMIP Keystore Encryption (BUI)

To configure encryption using KMIP, upload the key and certificates, and specify the KMIP server.

Before You Begin

Before setting up KMIP keystore on clustered controllers, perform the following procedures:

- For each cluster node, configure a private network resource that is able to reach the KMIP servers. This private network interface ensures that each cluster node can communicate with the KMIP server in case the data service network interfaces failover. For information about private resources, see Cluster Resource Management.
- Appropriately configure a route to each KMIP server for the private network link. For information about route configuration, see Configuring Network Routing.

> **Caution:**
>
> Failure to meet these prerequisites results in service interruption during takeover and failback operations.

1. From the **Configuration** menu, select **Settings**, then **Certificates**.
2. Upload the private key.

   a. Click the upload icon ⤒ to the left of **System**.

   b. Click **Browse**.

      Navigate to the location where your key and certificates are stored.

      For Oracle Key Vault, the key and certificates are contained in a `jar` file that you received from your Oracle Key Vault administrator.

   c. Select the PEM-format file that contains the private key.

      For Oracle Key Vault, this is the `key.pem` file.

   d. Click the **UPLOAD** button.

      A new row with a **Type** value of `key` appears in the **System** table.
3. Upload the system certificate.

   Select the PEM-format file that contains the client certificate for this system. For Oracle Key Vault, this is the `cert.pem` file.

   Follow the same steps that you used to upload the private key.

   The row that had a **Type** of `key` changes so that it now has a **Type** of `cert`. The certificate is matched to the existing key and combined into one object on the appliance. The **Subject**, **Issued By (CN)**, and **Expires** fields are populated. Click the information icon ℹ in the row to see the full values of these fields and more information.
4. Upload the trust anchor certificate.

   The CA certificate is the issuer of the client certificate.

   Select the PEM-format file that contains the CA certificate that issued the client certificate. For Oracle Key Vault, this is the `CA.pem` file.

   If your KMIP deployment uses multiple CA certificates, repeat this step to upload each certificate and set the services property to `kmip`. Multiple CA certificates might be necessary if the client certificate for Oracle ZFS Storage Appliance and the KMIP server certificate are issued from different CAs, or when there are intermediate CAs in the chain to the CA root certificate.

   a. Click the **Trusted** tab.

   b. Click the upload icon ⤒ to the left of **Trusted**.

   c. Click **Browse**.

   d. Select the CA certificate file.

   e. Click the **UPLOAD** button.

      A new row appears in the **Trusted** table.

   f. Click the edit icon ✎ in the new row.

   g. At the bottom of the **Certificate Details** dialog box, check the **kmip** box, and click **OK**.

      In the table row, the **Service** column now shows `kmip`.

5. From the **Shares** menu, select **Encryption**, then **KMIP**.

6. From the **Certificate** drop-down menu, select the certificate that you just uploaded.

7. Enter the hostname or IP address of the KMIP server.

   The recommended value to use is the KMIP server's hostname. See the discussion of hostname validation in Key Management Interoperability Protocol (KMIP) Keystore.

   If you specify an IP address, check with your KMIP server administrator regarding whether you need to include the port number.

8. Verify the **Match Hostname** and **Removing a key** options.

   By default, both of these options are enabled (checked). See Key Management Interoperability Protocol (KMIP) Keystore for information about what these options do.

9. Click **APPLY**.

   If you receive a warning that the server failed to validate in the certificate, click **Cancel** in the warning dialog box, and check whether you specified the correct server hostname in the **KMIP Server** field. If you specify an IP address for the KMIP server, and the CA-signed certificate subject common name only has a domain name, then host validation for the certificate fails.

   If you uncheck the **Match Hostname** option, host validation is not performed, and security is weaker.

10. Give the key a name.

    a. Click the add icon ⊕ to the left of **Keys**.

    b. Enter a **Keyname** and click **ADD**.

11. Optional: Identify pools, projects, and shares that are encrypted with this key.

    To identify which pools, projects, and shares are encrypted with this key, start the process of deleting the key, but cancel the key delete operation after you view the list of data that will be affected. When you initiate deleting the key, a warning is displayed that all data that is encrypted with this key will become inaccessible. The warning displays a list of all pools, projects, and shares that are encrypted with this key. Click the **Cancel** button and not the **OK** button to cancel deleting the key. For more details, see Deleting an Encryption Key (BUI).

Related Topics

- Key Management Interoperability Protocol (KMIP) Keystore
- Configuring LOCAL Keystore Encryption (BUI)
- Configuring OKM Keystore Encryption (BUI)

- Configuring KMIP Keystore Encryption (CLI)
- Creating an Encrypted Pool (BUI)
- Creating an Encrypted Project (BUI)

Configuring KMIP Keystore Encryption (CLI)

To configure encryption using KMIP, upload the key and certificates, and specify the KMIP server.

Before You Begin

Before setting up KMIP keystore on clustered controllers, perform the following procedures:

- For each cluster node, configure a private network resource that is able to reach the KMIP servers. This private network interface ensures that each cluster node can communicate with the KMIP server in case the data service network interfaces failover. For information about private resources, see Cluster Resource Management.
- Appropriately configure a route to each KMIP server for the private network link. For information about route configuration, see Configuring Network Routing.

> **Caution:**
>
> Failure to meet these prerequisites results in service interruption during takeover and failback operations.

1. Go to `configuration settings certificates system`.
2. Upload the private key.
   1. Copy the contents of the PEM-format file that contains the private key.

      For Oracle Key Vault, the key and certificates are contained in a `jar` file that you received from your Oracle Key Vault administrator. The private key file is the `key.pem` file.
   2. Enter the `import` command. Paste the key as prompted.

      ```
      hostname:configuration settings certificates system> import
      ("." to end)> -----BEGIN RSA PRIVATE KEY-----
      ...
      ("." to end)> -----END RSA PRIVATE KEY-----
      ("." to end)> .
      ```
   3. Enter the `list` command.

      The system certificates table has a new row with a `TYPE` value of `key`.

      ```
      hostname:configuration settings certificates system> list
      CERT     TYPE SUBJECT COMMON NAME       ISSUER COMMON NAME        NOT AFTER
      cert-001 key  RSA-2048
      ```
3. Upload the system certificate.

   The system certificate is the PEM-format file that contains the client certificate for this system. For Oracle Key Vault, this is the `cert.pem` file.

   Copy the contents of the certificate file, enter the `import` command, and paste the certificate.

```
hostname:configuration settings certificates system> import
("." to end)> -----BEGIN CERTIFICATE-----
...
("." to end)> -----END CERTIFICATE-----
("." to end)> .
```

4. Upload the trust anchor certificate.

   The CA certificate is the issuer of the client certificate. For Oracle Key Vault, this is the `CA.pem` file.

   If your KMIP deployment uses multiple CA certificates, repeat this step to upload each certificate, and set the services property to `kmip`. Multiple CA certificates might be necessary if the client certificate for Oracle ZFS Storage Appliance and the KMIP server certificate are issued from different CAs, or when there are intermediate CAs in the chain to the CA root certificate.

   a. Enter the `up trusted` command.

   b. Copy the contents of the certificate authority file, enter the `import` command, and paste the CA certificate.

      The trusted certificates table has a new row.

      ```
      hostname:configuration settings certificates trusted> list
      CERT     TYPE SUBJECT COMMON NAME       ISSUER COMMON NAME        NOT AFTER
      cert-001 cert CA                        CA                        2022-8-11
      ```

   c. Set the `services` property of the certificate to `kmip`.

      ```
      hostname:configuration settings certificates trusted> select cert-001
      hostname:configuration settings certificates cert-001> get services
                            services =
      hostname:configuration settings certificates cert-001> set services=kmip
                            services = kmip (uncommitted)
      hostname:configuration settings certificates cert-001> commit
      hostname:configuration settings certificates cert-001> done
      ```

5. Go to `shares encryption kmip`.

```
hostname:shares encryption kmip> get
                   server_list =
                   client_cert =
                    host_match = true
         destroy_key_on_remove = true
```

6. Set `client_cert` to the certificate that you just uploaded.

   By default, the `list` command lists KMIP servers. Use the `list certs` command to list available certificates.

```
hostname:shares encryption kmip> list certs
CERT     TYPE SUBJECT COMMON NAME   ISSUER COMMON NAME          NOT AFTER
cert-001 cert iogo7PmhIY            CA                          2023-1-25
```

   Set the value of `client_cert` by using either the `set` command or the `client_cert` command.

```
hostname:shares encryption kmip> set client_cert=cert-001
hostname:shares encryption kmip> client_cert cert-001
```

   For scripting, you need a persistent identifier for the certificate. Use the `client_cert` command with tab completion to list properties of certificates.

```
hostname:shares encryption kmip> client_cert tab
cert-001                        notafter
cert-002                        notbefore
comment                         sha1fingerprint
dns                             sha256fingerprint
ip                              subject_commonname
issuer_commonname               subject_countryname
issuer_countryname              subject_localityname
issuer_localityname             subject_organizationalunitname
issuer_organizationalunitname   subject_organizationname
issuer_organizationname         subject_stateorprovincename
issuer_stateorprovincename      type
key_bits                        uri
key_type                        uuid
```

Every certificate has a unique `subject_commonname`, so you can use the value of that property to set the `client_cert` property.

```
hostname:shares encryption kmip> client_cert subject_commonname=tab
                  ip-addr                         iogo7PmhIY
hostname:shares encryption kmip> client_cert subject_commonname=iogo7PmhIY
hostname:shares encryption kmip> get client_cert
                   client_cert = cert-001 (uncommitted)
```

7. Set `server_list` to the hostname or IP address of the KMIP server.

   The recommended value for this property is the KMIP server's hostname. See the discussion of hostname validation in Key Management Interoperability Protocol (KMIP) Keystore.

   If you specify an IP address, check with your KMIP server administrator regarding whether you need to include the port number.

   ```
   hostname:shares encryption kmip> set server_list=kmip-server-hostname-or-IP-
   address
                      server_list = kmip-server-hostname-or-IP-address (uncommitted)
   hostname:shares encryption kmip> commit
   ```

   If you receive a warning that the server failed to validate in the certificate, check whether you specified the correct server hostname in `server_list`. If you specify an IP address for `server_list`, and the CA-signed certificate subject common name only has a domain name, then host validation for the certificate fails.

   If you set the value of `host_match` to `false`, host validation is not performed, and security is weaker.

8. Verify the `host_match` and `destroy_key_on_remove` options.

   By default, both of these options are `true`. See Key Management Interoperability Protocol (KMIP) Keystore for information about what these options do.

9. Give the key a name.

   Create a key, and set the `keyname` property.

   ```
   hostname:shares encryption kmip> keys
   hostname:shares encryption keys> list
   NAME     CREATED               CIPHER KEYNAME
   hostname:shares encryption keys> create
   hostname:shares encryption kmip key-000 (uncommitted)> set keyname=atz-1-27-2021
                          keyname = atz-1-27-2021 (uncommitted)
   hostname:shares encryption kmip key-000 (uncommitted)> commit
   hostname:shares encryption keys> list
   ```

```
NAME     CREATED              CIPHER KEYNAME
key-000  2021-1-27 07:14:31   AES    atz-1-27-2021
```

10. Optional: Identify pools, projects, and shares that are encrypted with this key.

    To identify which pools, projects, and shares are encrypted with this key, start the process of deleting the key, but cancel the key delete operation after you view the list of data that will be affected. When you initiate deleting the key, a warning is displayed that all data that is encrypted with this key will become inaccessible. The warning displays a list of all pools, projects, and shares that are encrypted with this key. Enter `n` to cancel deleting the key.

    ```
    hostname:shares encryption keys> destroy key-000
    This key has the following dependents:
      pool-0/local/default/fs-enc
    Destroying this key will render the data inaccessible. Are you sure? (Y/N) n
    ```

Related Topics

- Key Management Interoperability Protocol (KMIP) Keystore
- Configuring LOCAL Keystore Encryption (CLI)
- Configuring OKM Keystore Encryption (CLI)
- Configuring KMIP Keystore Encryption (BUI)
- Creating an Encrypted Pool (CLI)
- Creating an Encrypted Project (CLI)

Creating an Encrypted Pool (BUI)

For more detailed information about creating a pool, see Creating a Storage Pool (BUI).

Pool keystore and key name values can be changed at any point after the pool has been created. However, you cannot add encryption information to a pool that was already created as unencrypted.

Any dataset that is created in an encrypted pool will also be encrypted. You cannot create unencrypted projects or shares on an encrypted pool.

Before You Begin

To create an encrypted storage pool, upgrade to software release OS8.8.0 or later and accept all deferred updates, including "Enable Pool Encryption." See Enable Pool Encryption Deferred Update in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The encryption key must be created before you can create an encrypted pool. See Data Encryption.

- Because the keystore must be configured before the pool is created, you cannot create an encrypted pool at initial system configuration or after factory reset.
- Before setting up replication for a share or project in a encrypted pool, ensure that the encryption key used at the source is also available at the target.

1. From the **Configuration** menu, select **Storage**.
2. Next to **Available Pools**, click the add icon ⊕ .
3. Type a name for the storage pool, and click **APPLY**.
4. Select the number of data drives for the storage pool for each disk shelf. You can also select available log, cache, and meta devices.

For more information about selecting data drives and meta devices, see Creating a Storage Pool (BUI).

5. Click **COMMIT**.

   The drives are allocated to the storage pool, and verified for presence and minimum functionality. If verification fails, click **ABORT**, fix the problem, and begin this procedure again. If you allocate a pool with missing or failed devices, you will not be able to add the missing or failed devices later.

6. On the **Choose Storage Profile** screen, select the data profile that meets your reliability, availability, serviceability, and performance goals.

   For a description of each profile, click on the data profile name, or see Data Profiles for Storage Pools.

7. If you allocated log, cache, or meta devices, select the appropriate profiles.

   For more information, see Creating a Storage Pool (BUI).

8. Set the encryption type, keystore, and key name.

   Use the fields **Encryption** and **Key** in the section **Optional Settings** at the bottom of the **Choose Storage Profile** screen.

   The **Encryption** field is disabled if no keystore is configured: The encryption key must be created before you create the pool. See Data Encryption.

   By default, **Encryption** is set to `Off` and **Key** is `disabled`. When you select a type in the **Encryption** field (see Understanding Encryption Key Values), then you must select a keystore and a key name in the **Key** field.

9. Click **COMMIT**.

   Once selected, the **Encryption** value is immutable. However, the **Key** values can be changed at any time. See Changing a Pool Encryption Key (BUI).

   All projects created under this pool are automatically encrypted with these encryption values, though the **Key** values can be changed. See Creating an Encrypted Project (BUI).

Related Topics

- Encryption Properties
- Managing Encryption Keys

Creating an Encrypted Pool (CLI)

For more detailed information about creating a pool, see Creating a Storage Pool (CLI).

Pool keystore and key name values can be changed at any point after the pool has been created. However, you cannot add encryption information to a pool that was already created as unencrypted.

Any dataset that is created in an encrypted pool will also be encrypted. You cannot create unencrypted projects or shares on an encrypted pool.

Before You Begin

To create an encrypted storage pool, upgrade to software release OS8.8.0 or later and accept all deferred updates, including "Enable Pool Encryption." See Enable Pool Encryption Deferred Update in *Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x*.

The encryption key must be created before you can create an encrypted pool. See Data Encryption.

- Because the keystore must be configured before the pool is created, you cannot create an encrypted pool at initial system configuration or after factory reset.
- Before setting up replication for a share or project in a encrypted pool, ensure that the encryption key used at the source is also available at the target.

1. Go to `configuration storage`.
2. Enter `config` and a name for the new storage pool.
3. Enter `show` to see the device information for the pool.
4. Enter `set` and the disk shelf or controller ID, and the number of data drives to use. You can also select available cache, meta, and log devices.

   For more information about selecting data drives and meta devices, see Creating a Storage Pool (CLI).
5. Enter `done`.
6. Enter `show` to display the profile.
7. If you allocated log devices to the pool, enter `set log_profile=` and set the log profile to either `log_mirror` or `log_stripe`. Use `log_mirror` if the pool contains an even number of log devices.
8. If you allocated meta devices to the pool, enter `set meta_profile=` and set the meta profile to either `meta_mirror` or `meta_stripe`.
9. Set the encryption type, keystore, and key name using properties `encryption`, `keystore`, and `keyname`.

   The properties `encryption`, `keystore`, and `keyname` are hidden and immutable if a keystore is not configured. The encryption key must be created before you create the pool. See Data Encryption.

   By default, `encryption` is set to `off`.

   ```
   hostname:configuration storage (pool1) config> ls
                PROFILE          CAPCTY NSPF DESCRIPTION
      profile = mirror           4.92G  no   Mirrored
                mirror3          4.92G  no   Triple mirrored
                stripe           14.8G  no   Striped

   encryption = off
      keyname =
     keystore =
   ```

   Set an encryption type (see Understanding Encryption Key Values), a keystore, and a key name. If `encryption` is not `off`, then `keystore` and `keyname` must be set.

   ```
   hostname:configuration storage (pool1) config> set encryption=aes-128-ccm
                       encryption = aes-128-ccm
   hostname:configuration storage (pool1) config> set keystore=LOCAL
                         keystore = LOCAL
   hostname:configuration storage (pool1) config> set keyname=MyKey
                          keyname = MyKey
   ```
10. Enter `done` to complete the task.

    ```
    hostname:configuration storage config (pool1)> done
    ```

    Check the values.

```
hostname:configuration storage (pool1)> get encryption keystore keyname
keystatus
                    encryption = aes-128-ccm
                      keystore = LOCAL
                       keyname = MyKey
                     keystatus = available
```

Once selected, the `encryption` value is immutable. However, the `keystore` and `keyname` values can be changed at any time. See Changing a Pool Encryption Key (CLI).

All projects created under this pool are automatically encrypted with these encryption values, though the `keystore` and `keyname` values can be changed. See Creating an Encrypted Project (CLI).

Related Topics

- Encryption Properties
- Managing Encryption Keys

Changing a Pool Encryption Key (BUI)

You can change the encryption key for a pool at any time, even while it is in use by client systems. Changing a pool encryption key changes the key for projects and shares that inherit the key from the pool.

1. From the **Configuration** menu, select **Storage**.
2. Select the pool for which you want to change the encryption key.
3. Change the pool encryption key.

   Click the key name or click the edit icon ✎ to the right of the **Key** field to display a dialog box for editing the keystore and key name used to encrypt the selected pool.
4. Click **APPLY**.

   All projects that are already inheriting encryption settings from this pool will also inherit the new encryption settings. See Creating an Encrypted Project (BUI).

Related Topics

- Deleting an Encryption Key (BUI)
- Encryption Properties

Changing a Pool Encryption Key (CLI)

You can change the encryption key for a pool at any time, even while it is in use by client systems. Changing a pool encryption key changes the key for projects and shares that inherit the key from the pool.

1. Go to `configuration storage`.
2. Select the pool for which you want to change the encryption key.

   If the pool that is selected is not the one you want, use the `set pool` command to select a different pool.
3. Change the pool encryption key.

   Set the values of the `keystore` and `keyname` properties.

4. Enter `commit`.

   All projects that are already inheriting encryption settings from this pool will also inherit the new encryption settings. See Creating an Encrypted Project (CLI).

Related Topics

- Deleting an Encryption Key (CLI)
- Encryption Properties

Creating an Encrypted Project (BUI)

A project inherits properties of the parent pool. A share inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A project that is created in an encrypted pool is automatically encrypted, and inherits the encryption settings from the pool. You can change the encryption key settings, but you cannot unencrypt the project and you cannot create an unencrypted project in an encrypted pool.

This procedure describes how to create an encrypted project in either an encrypted pool or an unencrypted pool.

Before You Begin

To use encryption, you must configure it first. See Data Encryption.

1. From the **Shares** menu, select **Projects**.
2. Click the add icon ⊕ next to **Projects** or in the expanded **Projects** panel. To expand the **Projects** panel, click its arrow icon.
3. In the **Create Project** dialog box, enter a name for the new project.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (.) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **`- . :`**
4. If the parent pool for this project is not encrypted, set an encryption type.

   If the parent pool for this project is not encrypted, the value of the **Encryption** field is `Off` by default. To create an encrypted project, set the **Encryption** field to a new value. See Understanding Encryption Key Values for descriptions of the values of the **Encryption** field.

   If the parent pool for this project is encrypted, the **Encryption** value is inherited, and you cannot change it.
5. Select a keystore and a key name.

   If the parent pool of the project is not encrypted, select a keystore and key name.

   If the parent pool of the project is encrypted, decide whether to inherit encryption settings from the pool. Any project or dataset that is created in an encrypted pool, including the default project, can be set to inherit the keystore and key name from the pool or can specify its own keystore and key name, and then the inherited key values will not be used.

   - To inherit the keystore and key name from the pool, check the **Inherit key** check box in the dialog box.

     If a project inherits encryption settings from the pool, then all of the project's datasets inherit encryption settings from the pool.

When you check the **Inherit key** check box, both **Key** selections (`keystore` and `key name`) are grayed out and not selectable. The key can only be modified from the parent pool.

- To set a different keystore and key name, uncheck the **Inherit key** check box. Select a keystore and key name.

6. Set other properties as appropriate for this project.

   Project properties are described in Project Properties.

7. Click **APPLY**.

   The new project is added to the **Projects** list.

   All shares created under this project are automatically encrypted with these encryption values, although the **Key** values can be changed. See Creating an Encrypted Filesystem or LUN (BUI).

Related Topics

- Encryption Properties
- Managing Encryption Keys
- Changing a Project Encryption Key (BUI)

Creating an Encrypted Project (CLI)

For more detailed information about creating a project, see Creating a Project (CLI).

A project inherits properties of the parent pool. A share inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A project that is created in an encrypted pool is automatically encrypted, and inherits the encryption settings from the pool. You can change the encryption key settings, but you cannot unencrypt the project and you cannot create an unencrypted project in an encrypted pool.

This procedure describes how to create an encrypted project in either an encrypted pool or an unencrypted pool.

Before You Begin

To use encryption, you must configure it first. See Data Encryption.

1. Go to `shares`.
2. Enter the `project` command, and a name for the project.
3. If the parent pool for this project is not encrypted, set an encryption type.

   If the parent pool for this project is not encrypted, the value of the `encryption` property is `off` by default. To create an encrypted project, set the `encryption` property to a new value. See Understanding Encryption Key Values for descriptions of the values of the `encryption` property.

   If the parent pool for this project is encrypted, the `encryption` property value is inherited and you cannot change it.

4. Set a keystore and a key name.

   If the parent pool for this project is not encrypted, set the `keystore` and `keyname` properties.

If the parent pool for this project is encrypted, the values of the `keystore` and `keyname` properties are inherited from the pool by default. You can use the inherited values or change them.

```
hostname:shares myproject (uncommitted)> get encryption keyname keystore
                    encryption = aes-128-ccm
                      keystore = LOCAL
                       keyname = MyKey
```

Use the `set` command to set new values for `keystore` and `keyname`, as shown in the following example:

```
hostname:shares myproject (uncommitted)> set keyname=NewKey
                       keyname = NewKey (uncommitted)
```

5. Use the `get` and `set` commands to set other properties as appropriate for this project.

   Project properties are described in Project Properties.

6. Enter `commit`.

   All shares created under this project are automatically encrypted with these encryption values, although the `keystore` and `keyname` values can be changed. See Creating an Encrypted Filesystem or LUN (CLI).

Related Topics

- Encryption Properties
- Managing Encryption Keys
- Changing a Project Encryption Key (CLI)

Changing a Project Encryption Key (BUI)

You can change the encryption key for a project at any time, even while it is in use by client systems. Changing a project encryption key changes the key for shares that inherit the key from the project.

If the parent pool of the project is encrypted, the project inherits encryption keys from the parent pool by default. You can change a key for a project without changing the key for the parent pool.

1. From the **Shares** menu, select **Projects**.
2. To find the project that you want to change, click **Show All**, **Local**, or **Replica**.
3. Move your cursor over the project that you want to change, and click the edit icon ✎ .
4. Click **General**.

   The project parameters are displayed.

5. If necessary, uncheck **Inherit from pool**.

   If **Inherit from pool** is checked, the key can only be modified from the parent pool.

6. To change the project **Encryption key**, select the keystore and the key name that you want to use.
7. Click **APPLY**.

   When you click **APPLY**, your changes are saved and the new key appears in the **Encryption key** area. All shares in this project inherit the new encryption settings. All shares that are already inheriting encryption settings from this project will also inherit the new encryption settings. See Creating an Encrypted Filesystem or LUN (BUI).

Related Topics

- Deleting an Encryption Key (BUI)
- Encryption Properties

Changing a Project Encryption Key (CLI)

You can change the encryption key for a project at any time, even while it is in use by client systems. Changing a project encryption key changes the key for shares that inherit the key from the project.

If the parent pool of the project is encrypted, the project inherits encryption keys from the parent pool by default. You can change a key for a project without changing the key for the parent pool.

1. Go to `shares`.
2. Select the project for which you want to change the encryption key.

   If the project that is selected is not the one you want, use the `select` *`project-name`* command to select a different project.
3. Change the project encryption key.

   Set the values of the `keystore` and `keyname` properties.
4. Enter `commit`.

   All shares that are already inheriting encryption settings from this project will also inherit the new encryption settings. See Creating an Encrypted Filesystem or LUN (CLI).

Related Topics

- Changing a Share Encryption Key (CLI)
- Deleting an Encryption Key (CLI)
- Encryption Properties

Creating an Encrypted Filesystem or LUN (BUI)

A filesystem or LUN that is created within a project inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A filesystem or LUN that is created within an encrypted project is automatically encrypted and inherits the encryption settings from the project. You can change the encryption key settings, whether the share is created in an unencrypted project or in an encrypted project.

An easy way to create a set of encrypted shares is to create them in an encrypted project.

Before You Begin

To use encryption, you must first configure a keystore and keys. See Data Encryption.

1. From the **Shares** menu, select **Shares**.
2. Select either **Filesystems** or **LUNs**.
3. Click the add icon ⊕ .

4. Complete the fields in the **Create Filesystem** or **Create LUN** dialog box.

   - For a filesystem, select a project and enter a name.
   - For a LUN, select a project, enter a name, and specify the volume size.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (.) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **- . :**

5. If the parent project for this filesystem or LUN is not encrypted, set an encryption type.

   If the parent project for this share is not encrypted, the value of the **Encryption** field is `Off` by default. To create an encrypted share, set the **Encryption** field to a new value. See Understanding Encryption Key Values for descriptions of the values of the **Encryption** field.

   If the parent project for this share is encrypted, the **Encryption** value is inherited, and you cannot change it.

6. Select a keystore and a key name.

   If the parent project of the share is not encrypted, select a keystore and key name.

   If the parent project of the share is encrypted, decide whether to inherit encryption settings from the project. Any share that is created in an encrypted project can be set to inherit the keystore and key name from the project or can specify its own keystore and key name, and then the inherited key values will not be used.

   - To inherit the keystore and key name from the project, check the **Inherit key** check box in the dialog box.

     Both **Key** selections (`keystore` and `key name`) are grayed out and not selectable. The key can only be modified from the parent project.

   - To set a different keystore and key name, uncheck the **Inherit key** check box.

     Select a keystore and key name.

7. Click **APPLY**.

Related Topics

- Encryption Properties
- Managing Encryption Keys
- Changing a Share Encryption Key (BUI)

Creating an Encrypted Filesystem or LUN (CLI)

A filesystem or LUN that is created within a project inherits properties of the parent project. For a list of properties that can be inherited, see Inherited Properties.

A filesystem or LUN that is created within an encrypted project is automatically encrypted and inherits the encryption settings from the project. You can change the encryption key settings, whether the share is created in an unencrypted project or in an encrypted project.

An easy way to create a set of encrypted shares is to create them in an encrypted project.

Before You Begin

To use encryption, you must configure it first. See Data Encryption.

1. Go to `shares`.

2. Select the project.

   If the project that is selected is not the one you want, use the `select` *`project-name`* command to select a different project.

3. Create the filesystem or LUN.

   Enter `filesystem` *`filesystem-name`* or `lun` *`lun-name`*.

   The name must be 1 to 64 characters in length. The name cannot begin with a period (.) and cannot include spaces. Allowable characters are alphanumeric characters and special characters _ **`- . :`**

   The following example creates a filesystem named `fs-1` in the `default` project.

```
hostname:shares default> filesystem fs-1
hostname:shares default/fs-1 (uncommitted)>
```

4. If you are creating a LUN, enter `set volsize=` and the volume size.

```
hostname:shares default/lun1 (uncommitted)> set volsize=2G
                       volsize = 2G (uncommitted)
```

5. If the parent project for this filesystem or LUN is not encrypted, set an encryption type.

   If the parent project for this share is not encrypted, the value of the `encryption` property is `off` by default. To create an encrypted share, set the `encryption` property to a new value. See Understanding Encryption Key Values for descriptions of the values of the `encryption` property.

   If the parent project for this share is encrypted, the `encryption` value is inherited, and you cannot change it.

6. Set a keystore and a key name.

   If the parent project of the share is not encrypted, set the `keystore` and `keyname` properties.

   If the parent project for this share is encrypted, the values of the `keystore` and `keyname` properties are inherited from the project by default. You can use the inherited values or you can change them as shown in the following example:

```
hostname:shares default/fs-1 (uncommitted)> set keyname=MyKey
                       keyname = MyKey (uncommitted)
```

7. Use the `get` and `set` commands to set other properties as appropriate for this filesystem or LUN.

   Share properties are described in Filesystem Properties and LUN Properties.

8. Enter `commit`.

```
hostname:shares default/fs-1 (uncommitted)> commit
```

Related Topics

- Encryption Properties
- Managing Encryption Keys
- Changing a Share Encryption Key (CLI)

Changing a Share Encryption Key (BUI)

You can change the encryption key for a share at any time, even while it is in use by client systems. By default, shares inherit encryption keys from the parent project. You can change a key for a share without changing the key for the parent project.

1. From the **Shares** menu, select **Shares**.
2. Select **Filesystems** or **LUNs**.
3. To find the share that you want to change, click **Show All**, **Local**, or **Replica**.
4. Move your cursor over the share that you want to change, and click the edit icon ✎ .

   The share properties are displayed.
5. If necessary, uncheck **Inherit from project**.

   If **Inherit from project** is checked, the key can only be modified from the parent project.
6. To change the encryption key, select the keystore and the key that you want to use.
7. Click **APPLY**.

   When you click **APPLY**, your changes are saved and the new key appears in the **Encryption key** area.

Related Topics

- Changing a Project Encryption Key (BUI)
- Deleting an Encryption Key (BUI)
- Encryption Properties

Changing a Share Encryption Key (CLI)

You can change the encryption key for a share at any time, even while it is in use by client systems. By default, shares inherit encryption keys from the parent project. You can change a key for a share without changing the key for the parent project.

1. Select the filesystem or LUN for which you want to change the encryption key.
   a. Go to `shares`.
   b. Select the project.

      If the project that is selected is not the one you want, use the `select` *`project-name`* command to select a different project.
   c. Select the filesystem or LUN.

      Enter `select` *`filesystem-name`* or `select` *`lun-name`*.
2. Change the project encryption key.

   Set the values of the `keystore` and `keyname` properties.
3. Enter `commit`.

Related Topics

- Changing a Project Encryption Key (CLI)
- Deleting an Encryption Key (CLI)

- Encryption Properties

Backing Up a LOCAL Key (BUI)

Use the following procedure to retrieve the information for a single LOCAL key in order to back it up.

1. From the **Shares** menu, select **Encryption**, then **Local**.
2. Click on the key you want to back up.

   A dialog box appears with the **Keyname** and **Key** value.

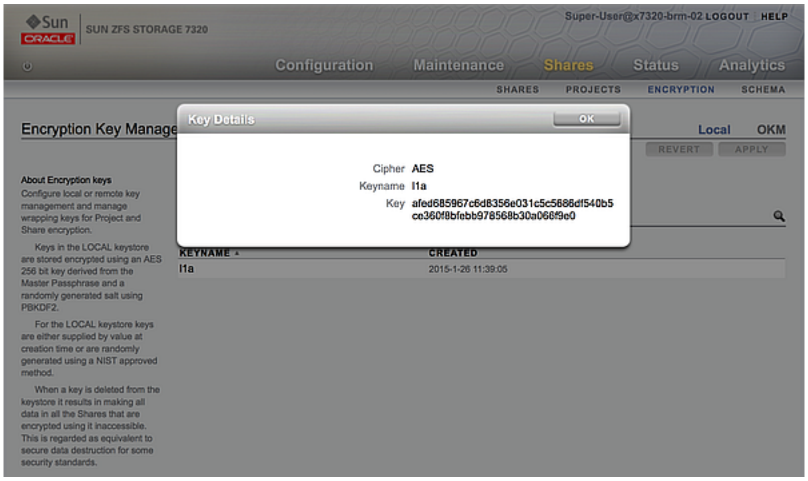

Cipher AES
Keyname l1a
Key afed685967c6d8356e031c5c5686df540b5
ce360f8bfebb978568b30a066f9e0

KEYNAME	CREATED
l1a	2015-1-26 11:39:05

About Encryption keys
Configure local or remote key management and manage wrapping keys for Project and Share encryption.

Keys in the LOCAL keystore are stored encrypted using an AES 256 bit key derived from the Master Passphrase and a randomly generated salt using PBKDF2.

For the LOCAL keystore keys are either supplied by value at creation time or are randomly generated using a NIST approved method.

When a key is deleted from the keystore it results in making all data in all the Shares that are encrypted using it inaccessible. This is regarded as equivalent to secure data destruction for some security standards.

3. Using any method, record this information in a backup location of your choosing, and then click **OK**.

Backing Up a LOCAL Key (CLI)

Use the following procedure to retrieve the information for a single LOCAL key in order to back it up.

1. Select the key.

```
 hostname:shares encryption local keys> select keyname=Mykey
```

2. Get the key value.

```
hostname:shares encryption local key-005> get key
     key = d6a5b801ffb93fcb19ef70a11d662d8092f243c5d4ccd0cd34264b15dd0b7739
```

3. Using any method, record this information in a backup location of your choosing.

Deleting an Encryption Key (BUI)

Deleting an encryption key is a fast and effective way to make large amounts of data inaccessible. Keys can be deleted even if they are in use. If the key is in use, a warning is given and confirmation is required. All shares, projects, or pools that use that key are unshared and can no longer be accessed by clients.

If you might use a LOCAL key again to access its associated shares, back up the key name and value before deleting the key as described in Backing Up a LOCAL Key (BUI). Then you can later perform a restore procedure as described in Restoring a LOCAL Key (BUI).

When an encryption key that is in use by a pool, project, or share is deleted, all affected pools, projects, and shares are listed as dependents for the key in the **Key Destroy** dialog box. When the key is deleted, the **Key Status** value changes to `unavailable`, and a warning indicator ● is displayed to the right of the lock icon 🔒 for the affected pool or share.

Use the following procedure to delete an encryption key.

1. From the **Shares** menu, select **Encryption**.
2. Select the appropriate keystore tab.
3. Move your cursor over the key that you want to delete, and click the delete icon 🗑 .

   An alert is displayed that warns you that all shares that are using this key will be unmounted and unshared. If you delete this key, all data in the shares that are encrypted using this key will be permanently and irrecoverably inaccessible. Then the alert lists the pools and shares that depend on this key.
4. To delete the key, click **OK**. To keep the key, click **CANCEL**.

   When a key is deleted, all of the data in all of the pools and shares that use the key becomes inaccessible. This is equivalent to secure data destruction and is permanent and irrevocable unless you have prepared for key restoration by backing up the key. For more information about key backup and restoration, see Backing Up a LOCAL Key (BUI) and Restoring a LOCAL Key (BUI).

Related Topics

- Changing a Share Encryption Key (BUI)
- Managing Encryption Keys
- Encryption Key Life Cycle

Deleting an Encryption Key (CLI)

Deleting an encryption key is a fast and effective way to make large amounts of data inaccessible. Keys can be deleted even if they are in use. If the key is in use, a warning is given and confirmation is required. All shares, projects, or pools that use that key are unshared and can no longer be accessed by clients.

If you might use a LOCAL key again to access its associated shares, back up the key name and value before deleting the key. Then you can later perform a restore procedure as described in Restoring a LOCAL Key (CLI).

When an encryption key that is in use by a pool, project, or share is deleted, all affected pools, projects, and shares are listed as dependents for the key. When the key is deleted, the `keystatus` property value changes to `unavailable`.

Use the following procedure to delete an encryption key.

1. Go to `shares encryption`.
2. Go to the appropriate keystore configuration.
3. Delete the key.

   Use the `destroy keyname` command to delete a key.

   ```
   hostname:shares encryption local local_keys> destroy keyname=MyKey

   This key has the following dependent shares:
   ```

```
    pool-1
    pool-1/local/default
    pool-1/local/default/fs-1

Destroying this key will render the data inaccessible. Are you sure? (Y/N)
```

Read the warning and confirm that you want to delete the key.

When a key is deleted, all of the data in all of the pools and shares that use the key becomes inaccessible. This is equivalent to secure data destruction and is permanent and irrevocable unless you have prepared for key restoration by backing up the key. For more information about key backup and restoration, see Backing Up a LOCAL Key (CLI) and Restoring a LOCAL Key (CLI).

4. Verify that a share is no longer accessible by using that key.

   - The value of the `keystatus` property changes to `unavailable`.
   - The `keystatus` property is marked as a critical property.
   - The missing key is treated as an error.

```
hostname:> shares select default select fs-1
hostname:shares default/fs-1> get encryption keystore keyname keystatus

                   encryption = aes-128-ccm (inherited)
                     keystore = LOCAL (inherited)
                      keyname = MyKey (inherited)
                    keystatus = unavailable

Errors:
            key_unavailable
```

5. To list dependents, use the following CLI commands:

```
hostname:shares (pool-1) encryption local keys> select keyname=1 hostname:shares
       (pool-010) encryption local key-002> list

Properties:
                        cipher = AES
                       keyname = 1

hostname:shares (pool-010) encryption local key-002> list dependents DEPENDENTS
        pool-010/local/default/a hostname:shares (pool-010) encryption local key-002>
```

Related Topics

- Changing a Share Encryption Key (CLI)
- Backing Up a LOCAL Key (CLI)
- Restoring a LOCAL Key (CLI)

Restoring a LOCAL Key (BUI)

To restore a LOCAL key that was deleted, create a new LOCAL key with the same key name and value as the deleted key. You must have first recorded, or backed up, this information before the key was deleted. The backup procedure is described in Backing Up a LOCAL Key (BUI). Although deleting a LOCAL key renders shares inaccessible, the shares can be made accessible again by recreating the LOCAL key. For information about restoring keys stored in the OKM keystore, refer to the Oracle Key Manager documentation on the Oracle Help Center.

Use the following procedure to restore a backed up LOCAL key.

Note:

If the key name is in use with a different key value for existing shares, change the key used for those shares before restoring the original LOCAL key. For more information, see Changing a Share Encryption Key (BUI).

1. Retrieve the key name and value for the LOCAL key from your backup location.
2. From the **Shares** menu, select **Encryption**, then **Local**, and click the add icon ⊕ .
3. Enter the same key name as in the backup.
4. Uncheck **Generate key automatically** and set the key value based on the backup.
5. Save the restored key by clicking **ADD**.

   If the key name is used with existing shares, a dialog box appears. To overwrite the key value in the existing shares, click **OK**. Click **CANCEL** to not add the new key. You can then change the key used for those shares before repeating this procedure and restoring the original key. For more information, see Changing a Share Encryption Key (BUI).

Related Topics

- Changing a Share Encryption Key (BUI)
- Backing Up a LOCAL Key (BUI)
- Deleting an Encryption Key (BUI)

Restoring a LOCAL Key (CLI)

To restore a LOCAL key that was deleted, create a new LOCAL key with the same key name and value as the deleted key. You must have first recorded, or backed up, this information before the key was deleted. The backup procedure is described in Backing Up a LOCAL Key (CLI). Although deleting a LOCAL key renders shares inaccessible, the shares can be made accessible again by recreating the LOCAL key. For information about restoring keys stored in the OKM keystore, refer to the Oracle Key Manager documentation on the Oracle Help Center.

Use the following procedure to restore a backed up LOCAL key.

Note:

If the key name is in use with a different key value for existing shares, change the key used for those shares before restoring the original LOCAL key. For more information, see Changing a Share Encryption Key (CLI).

1. Retrieve the key name and value for the LOCAL key from your backup location.
2. Create a key in the LOCAL keystore.

```
hostname:shares encryption local keys> create
```

3. Name the key based on the backup.

```
hostname:shares encryption local key-005 (uncommitted)> set keyname=Mykey
     keyname = Mykey (uncommitted)
```

4. Set the key value based on the backup.

```
hostname:shares encryption local key-005 (uncommitted)> set key=key-value
     key = key-value (uncommitted)
```

5. Save the key.

```
hostname:shares encryption local key-005 (uncommitted)> commit
```

   If the key name is used with existing shares, you will be alerted:

```
Existing shares reference the key Mykey from the LOCAL keystore. Are you sure? (Y/N)
```

   To overwrite the key value in the existing shares, type `Y`. Type `N` to not add the new key. You can then change the key used for those shares before repeating this procedure and restoring the original key. For more information, see Changing a Share Encryption Key (CLI).

Related Topics

- Changing a Share Encryption Key (CLI)
- Backing Up a LOCAL Key (CLI)
- Deleting an Encryption Key (CLI)

Encryption Properties

The following list shows the encryption properties available for creating keys, managing keys, and creating encrypted pools, projects, and shares.

- **LOCAL Key Management Properties**
  - **Master Passphrase -** The master passphrase is used to generate an AES key for encrypting the keys stored in the LOCAL keystore. The PKCS#5 PBKDF algorithm is used to generate the key, and the key is randomly generated and managed by the system.
- **LOCAL Key Creation Properties**
  - **Keyname -** Name to identify the key.
  - **Key -** Hex-encoded raw 256-bit key, stored in an encrypted form.
- **OKM Key Management Properties** (supplied by your OKM administrator)
  - **Key Manager Server -** IP address of your OKM server.
  - **User Agent ID** - Agent ID.
  - **Registration PIN** - Registration PIN.
- **OKM Key Creation Properties**
  - **Keyname** - Name to identify the key.
- **KMIP Key Management Properties**
  - **KMIP Server** - Hostname or IP address of a KMIP server. This property can have multiple values.
  - **Certificate** - Certificate that you uploaded to the appliance from files provided by your KMIP server administrator.
- **KMIP Key Creation Properties**
  - **Keyname** - Name to identify the key.

- **KMIP Options** - See Key Management Interoperability Protocol (KMIP) Keystore for more information.
  - Validate the server hostname against the server's identity in the server certificate.
  - Destroy or preserve a key on the KMIP server when that key is deleted key on the appliance.
- **Pool, Project, and Shares Encryption Properties**
  - **Encryption** - AES encryption type and key length. For more information, see Understanding Encryption Key Values.
  - **Keystore** - LOCAL, OKM, or KMIP.
  - **Key** - The name of a specific LOCAL, OKM, or KMIP key.
  - **Key Last Change -** The date that the key was last changed.
  - **Key Status** - If the value of this property is `unavailable`, then the key has been deleted.

Related Topics

- Data Encryption Workflow
- Managing Encryption Keys
- Performance Impact of Encryption
- Encryption Key Life Cycle

Managing Encryption Keys

Oracle ZFS Storage Appliance includes a built-in LOCAL keystore and the ability to connect to OKM and KMIP keystores. Each encrypted pool, project, or share requires a wrapping key from a keystore. The data encryption keys are managed by the storage appliance and are stored persistently encrypted by the wrapping key from the keystore.

Oracle Key Manager (OKM) Keystore

Oracle Key Manager (OKM) is a comprehensive key management system (KMS) that addresses the rapidly growing enterprise need for storage-based data encryption. Developed to comply with open standards, this feature provides the capacity, scalability, and interoperability to manage encryption keys centrally over widely distributed and heterogeneous storage infrastructures.

OKM meets the unique challenges of storage key management, including:

- **Long-term key retention** - OKM ensures that archive data is always available, and it securely retains encryption keys for the full data life cycle.
- **Interoperability** - OKM provides the interoperability needed to support a diverse range of storage devices attached to mainframe or open systems under a single storage key management service.
- **High availability** - With active N-node clustering, dynamic load balancing, and automated failover, OKM provides high availability, whether the appliances are sited together or distributed around the world.
- **High capacity** - OKM manages large numbers of storage devices and even more storage keys. A single clustered appliance can provide key management services for thousands of storage devices and millions of storage keys.

- **Flexible key configuration** - Per OKM cluster, keys can be generated automatically or created individually for a LOCAL or OKM keystore. Security administrators are responsible for providing the key names which, when combined with the keystore, associate a given wrapping key with a pool, project, or share.

> **Note:**
>
> If the Oracle ZFS Storage Appliance system is clustered, do not use the "one time passphrase" setting when creating the OKM server agent. Otherwise, registration on the other cluster node will fail, and keys will not be available on failover.

Key Management Interoperability Protocol (KMIP) Keystore

The Key Management Interoperability Protocol (KMIP) keystore is used in conjunction with KMIP-compliant servers, including Oracle Key Vault. Oracle Key Vault is a software appliance that is installed on a dedicated server and that supports the OASIS KMIP standard.

When multiple KMIP servers are listed, Oracle ZFS Storage Appliance will failover to alternate servers if the current server is not responding. Each configured KMIP server must present the same set of keys to the appliance, and must accept the same certificate presented by the appliance for client authentication.

The set of KMIP servers and the client certificate can be changed without removing the keys from the keystore.

Match Hostname

When this option is enabled, the system attempts to verify that the specified KMIP server matches the host specified in the peer server certificate.

KMIP allows you to use either the hostname or the IP address to specify a KMIP server. If you specify an IP address for the KMIP server, and the CA-signed certificate subject common name only has a domain name, then host validation for the certificate fails. If the **Match Hostname** BUI option is disabled or the `host_match` CLI property is set to `false`, host validation is not performed.

For stronger security, perform the host validation. Use the hostname to specify the KMIP server, and enable the host validation option.

Destroy or Preserve a Key on the Server

KMIP has the option to destroy or preserve a key on the KMIP server when that key is deleted on the appliance. When the option is enabled, a key that is deleted from the list of keys that are known to the appliance is also destroyed on the key server. When the option is disabled, keys remain on the key server after they are deleted from the list of keys on the appliance.

One example of when you want to keep keys on the key server after they are deleted from the appliance is when multiple, separate appliances are configured in Oracle Key Vault to see the same set of keys. If you delete a key on one appliance and the key is deleted from the key server, the key remains on the list on the other appliances and you cannot delete the key because the key is not found and appears to be already deleted.

Another example of when you want to keep keys on the key server after they are deleted from the appliance is when you are using replication to repurpose the appliance. You move (replicate) all the source off the appliance and you want to encrypt that moved source with the same keys. When the replication process cleans the original source appliance and deletes the key, the shares on the replica will not be able to be encrypted with that key if the key is deleted from the key server.

Maintaining Keys

Shares, projects, and pools that use OKM or KMIP keys that are in a deactivated state remain accessible. To prevent an OKM or KMIP key from being used, you must explicitly delete the key.

To ensure encrypted shares, projects, and pools are accessible, back up your appliance configurations and LOCAL keystore key values. If a key becomes unavailable, any shares, projects, or pools that use that key become inaccessible.

- If a pool key is unavailable, new projects cannot be created in that pool.
- If a project key is unavailable, new shares cannot be created in that project.

Keys can become unavailable in the following ways:

- Keys are deleted
- Rollback to a release that does not support encryption
- Rollback to a release where the keys are not configured
- Factory reset
- OKM or KMIP server is not available

Understanding Encryption Key Values

The following table shows the BUI and CLI encryption key values and descriptions. It also indicates if the encryption type works with deduplication.

Table 8-1 Encryption Key Values

BUI Value	CLI Value	Description
Off	`off`	Share/project/pool is not encrypted
AES-128-CCM	`aes-128-ccm`	Lowest CPU impact encryption, Dedupable
AES-192-CCM	`aes-192-ccm`	Dedupable
AES-256-CCM	`aes-256-ccm`	Dedupable
AES-128-GCM	`aes-128-gcm`	NIST SP800-38D recommended, Not-Dedupable
AES-192-GCM	`aes-192-gcm`	NIST SP800-38D recommended, Not-Dedupable
AES-256-GCM	`aes-256-gcm`	Highest CPU impact encryption, NIST SP800-38D recommended, Not-Dedupable

Performance Impact of Encryption

Using encryption with shares can have CPU performance impacts, as follows:

- The AES-128-CCM mode has the lowest CPU performance impact and is recommended for all workloads where there are no LOCAL security requirements.

- When encrypted data is read, it is stored decrypted and decompressed in DRAM. For read-dominant workloads that can be serviced read-dominant from the DRAM cache, the impact of decrypting the data is minimal.
- When SSD cache devices are used, data blocks evicted out of DRAM to the cache are compressed and encrypted, and must be decrypted and decompressed when retrieved back into DRAM.
- For workloads that are write-dominant and use larger block sizes, especially 128 kilobytes and 1 megabyte, there can be a significant CPU impact resulting in lower throughput. This is particularly likely if the filesystem record size or LUN volume block size is larger than the application block size.

Related Topics

- Data Encryption Workflow
- Encryption Properties
- Managing Encryption Keys
- Encryption Key Life Cycle

Encryption Key Life Cycle

The encryption key life cycle is flexible because you can change keys at any time without taking data services offline. When a key is deleted from the keystore, all the shares that use that key are unmounted and their data becomes inaccessible.

- Backup of keys in the LOCAL keystore is included as part of the System Configuration Backup. For the LOCAL keystore, it is also possible to supply the key by value at creation time to allow it to be escrowed in an external system, which provides an alternative per-key backup/restore capability.
- Backing up keys in the OKM keystore should be performed using the OKM backup services.
- Backing up keys in the KMIP keystore should be performed by the appropriate service on the KMIP server.

Related Topics

- Data Encryption Workflow
- Encryption Properties
- Managing Encryption Keys
- Performance Impact of Encryption

Backing Up, Replicating, and Restoring Encrypted Projects and Shares

Use one of the following methods for backing up and restoring encrypted projects and shares:

- **NDMP** - See NDMP Configuration.
- **Remote replication** - See Remote Replication.

For remote replication of encrypted projects or shares, both the source and target must support encryption, and meet the following requirements:

- Software release 2013.1.3.0 or later.
- Encryption wrapping keys used by the project or share are available.
- OKM or KMIP key name must be identical in the keystore on both replication source and replication targets.
- OKM Agent ID or KMIP host name must be unique on the replication source and target replication appliances. Replication peer appliances cannot use the same agent or host.
- OKM agents or KMIP servers for the replication peers should be configured on the OKM or KMIP server to see the same key groups.

The replication will fail if you attempt to replicate an encrypted project or share and the target does not support encryption. If the wrapping key is not available on the source or target system, or the target software is earlier than software release 2013.1.3.0, an alert is raised. Review the alerts on both the source and target to determine the reason for the replication failure.

Related Topics

- Data Encryption Workflow
- Managing Encryption Keys
- Encryption Key Life Cycle
- *Oracle ZFS Storage Appliance: Remote Replication Compatibility* (Doc ID 1958039.1) on https://support.oracle.com/

Inherited Property Settings after Restoration from Backup

In this description, “original” is the project or share that is backed up, prior to restoration.

If the original project or share inherits properties from its parent pool or project, then the restored project or share inherits properties from its target parent. If the original parent has different property settings from those of the target parent, then the inherited property settings of the project or share are changed when the project or share is restored.

If you want the restored project or share to have the same property settings as the original project or share, configure those settings explicitly on the original project or share, instead of having the original project or share inherit the settings.

For example, if an encrypted share inherits encryption property settings from its parent project, specify those encryption property settings explicitly on the share before backing up and restoring the share to a different project that has different encryption settings.

9

Maintenance Workflows

A workflow is a CLI script that is uploaded to and managed by Oracle ZFS Storage Appliance by itself. Workflows can be parameterized and executed in a first-class fashion from either the browser interface or the command line interface. Workflows may also be optionally executed as alert or at a designated time. As such, workflows allow for the appliance to be *extended* in ways that capture specific policies and procedures, and can be used (for example) to formally encode best practices for a particular organization or application.

To use workflows, use the following sections:

- Understanding Workflows
- Understanding Workflow Parameters
- Constrained Workflow Parameters
- Optional Workflow Parameters
- Workflow Error Handling
- Workflow Input Validation
- Workflow Execution Auditing and Reporting
- Understanding Workflow Versioning
- Using Workflows for Alert Actions
- Creating and Posting Custom Alerts from Within a Workflow
- Using Scheduled Workflows
- Using a Scheduled Workflow
- Coding Workflow Schedules
- Creating a Worksheet Based on a Specified Drive Type
- Uploading and Executing Workflows using the BUI
- Downloading Workflows using the CLI
- Listing Workflows using the CLI
- Executing Workflows using the CLI
- Auditing Workflows using the CLI

Understanding Workflows

A workflow is embodied in a valid ECMAScript file that contains a single global variable: `workflow`. The `workflow` object must contain at least three members:

Table 9-1 Required `workflow` Object Members

Required Member	Type	Description
`name`	String	Name of the workflow

Table 9-1 (Cont.) Required workflow Object Members

Required Member	Type	Description
description	String	Description of the workflow
execute	Function	Function that executes the workflow

Example 9-1 Hello World Workflow

This example shows a simple workflow.

```
var workflow = {
       name: 'Hello world',
       description: 'Bids a greeting to the world',
       execute: function () { return ('hello world!') }
};
```

Uploading this workflow results in a new workflow named "Hello world". Executing this workflow results in the output "hello world!"

Example 9-2 Using the Workflow Run Function to Return CPU Utilization

Workflows execute asynchronously in the appliance shell, running (by default) as the user that is executing the workflow. As such, workflows have at their disposal the appliance scripting facility (see Working with CLI Scripting), and can interact with the appliance in the same way as any other instance of the appliance shell. For example, workflows can execute commands, parse output, and modify state. This more complex example uses the `run` function to return the current CPU utilization.

```
var workflow = {
       name: 'CPU utilization',
       description: 'Displays the current CPU utilization',
       execute: function () {
               run('analytics datasets select name=cpu.utilization');
               cpu = run('csv 1').split('\n')[1].split(',');
               return ('At ' + cpu[0] + ', utilization is ' + cpu[1] + '%');
       }
};
```

Understanding Workflow Parameters

Workflows that do not operate on input have limited scope; many workflows need to be parameterized to be useful. This is done by adding a `parameters` member to the global `workflow` object. The `parameters` member is in turn an object that is expected to have a member for each parameter. Each `parameters` member must have the following members:

Table 9-2 Required Workflow Parameters Members

Required Member	Type	Description
label	String	Label to adorn input of workflow parameter
type	String	Type of workflow parameter

The `type` member must be set to one of these types:

Table 9-3 Workflow Member Type Names

Type Name	Description
`Boolean`	A boolean value
`ChooseOne`	One of a number of specified values
`EmailAddress`	An e-mail address
`File`	A file to be transferred to the appliance
`Host`	A valid host, as either a name or dotted decimal
`HostName`	A valid hostname
`HostPort`	A valid, available port
`Integer`	An integer
`NetAddress`	A network address
`NodeName`	A name of a network node
`NonNegativeInteger`	An integer that is greater than or equal to zero
`Number`	Any number, including floating point
`Password`	A password
`Permissions`	POSIX permissions
`Port`	A port number
`Size`	A size
`String`	A string
`StringList`	A list of strings

Example 9-3 Workflow Using Two Parameters

Based on the specified types, an appropriate input form will be generated upon execution of the workflow. For example, here is a workflow that has two parameters: the name of a business unit (to be used as a project), and the name of a share (to be used as the share name).

```
var workflow = {
       name: 'New share',
       description: 'Creates a new share in a business unit',
       parameters: {
               name: {
                       label: 'Name of new share',
                       type: 'String'
               },
               unit: {
                       label: 'Business unit',
                       type: 'String'
               }
       },
       execute: function (params) {
               run('shares select ' + params.unit);
               run('filesystem ' + params.name);
               run('commit');
               return ('Created new share "' + params.name + '"');
       }
};
```

If you upload this workflow and execute it, you will be prompted with a dialog box to fill in the name of the share and the business unit. When the share has been created, a message will be generated indicating as much.

Constrained Workflow Parameters

For some parameters, you might not want to allow an arbitrary string, but would rather limit input to one of a small number of alternatives. These parameters should be specified to be of type `ChooseOne`, and the object containing the parameter must have two additional members:

Table 9-4 Constrained Parameters Required Members

Required Member	Type	Description
`options`	Array	An array of strings that specifies the valid options
`optionlabels`	Array	An array of strings that specifies the labels associated with the options specified in `options`

Example 9-4 Using the Workflow `ChooseOne` Parameter

Using the `ChooseOne` parameter type, we can enhance the previous example to limit the business unit to be one of a small number of predefined values:

```
var workflow = {
    name: 'Create share',
    description: 'Creates a new share in a business unit',
    parameters: {
        name: {
            label: 'Name of new share',
            type: 'String'
        },
        unit: {
            label: 'Business unit',
            type: 'ChooseOne',
            options: [ 'development', 'finance', 'qa', 'sales' ],
            optionlabels: [ 'Development', 'Finance',
                'Quality Assurance', 'Sales/Administrative' ],
        }
    },
    execute: function (params) {
        run('shares select ' + params.unit);
        run('filesystem ' + params.name);
        run('commit');
        return ('Created new share "' + params.name + '"');
    }
};
```

When this workflow is executed, the `unit` parameter will not be entered manually; it will be selected from the specified list of possible options.

Optional Workflow Parameters

Some parameters may be considered *optional* in that the UI should not mandate that these parameters are set to any value to allow execution of the workflow. Such a parameter is denoted via the `optional` field of the `parameters` member:

Table 9-5 Required Members for Optional Parameters

Optional Member	Type	Description
`optional`	Boolean	If set to `true`, denotes that the parameter need not be set; the UI may allow the workflow to be executed without a value being specified for the parameter

If a parameter is optional and is unset, its member in the parameters object passed to the `execute` function will be set to `undefined`.

Workflow Error Handling

If, in the course of executing a workflow, an error is encountered, an exception will be thrown. If the exception is not caught by the workflow itself (or if the workflow throws an exception that is not otherwise caught), the workflow will fail, and the information regarding the exception will be displayed to the user. To properly handle errors, exceptions should be caught and processed. For example, in the previous example, an attempt to create a share in a non-existent project results in an uncaught exception.

Example 9-5 Workflow Error Handling

This example could be modified to catch the offending error, and create the project if it does not exist:

```
var workflow = {
    name: 'Create share',
    description: 'Creates a new share in a business unit',
    parameters: {
        name: {
            label: 'Name of new share',
            type: 'String'
        },
        unit: {
            label: 'Business unit',
            type: 'ChooseOne',
            options: [ 'development', 'finance', 'qa', 'sales' ],
            optionlabels: [ 'Development', 'Finance',
                'Quality Assurance', 'Sales/Administrative' ],
        }
    },
    execute: function (params) {
        try {
            run('shares select ' + params.unit);
        } catch (err) {
            if (err.code != EAKSH_ENTITY_BADSELECT)
                throw (err);

            /*
             * We haven't yet created a project that corresponds to
             * this business unit; create it now.
             */
            run('shares project ' + params.unit);
            run('commit');
            run('shares select ' + params.unit);
        }

        run('filesystem ' + params.name);
        run('commit');
        return ('Created new share "' + params.name + '"');
    }
};
```

Workflow Input Validation

Workflows may optionally validate their input by adding a `validate` member that takes as a parameter an object that contains the workflow parameters as members. The `validate` function should return an object where each member is named with the parameter that failed validation, and each member's value is the validation failure message to be displayed to the user.

Example 9-6 Workflow Input Validation

This extends our example to give a precise error if the user attempts to create an extant share:

```
var workflow = {
    name: 'Create share',
    description: 'Creates a new share in a business unit',
    parameters: {
        name: {
            label: 'Name of new share',
            type: 'String'
        },
        unit: {
            label: 'Business unit',
            type: 'ChooseOne',
            options: [ 'development', 'finance', 'qa', 'sales' ],
            optionlabels: [ 'Development', 'Finance',
                'Quality Assurance', 'Sales/Administrative' ],
        }
    },
    validate: function (params) {
        try {
            run('shares select ' + params.unit);
            run('select ' + params.name);
        } catch (err) {
            if (err.code == EAKSH_ENTITY_BADSELECT)
                return;
        }

        return ({ name: 'share already exists' });
    },
    execute: function (params) {
        try {
            run('shares select ' + params.unit);
        } catch (err) {
            if (err.code != EAKSH_ENTITY_BADSELECT)
                throw (err);

            /*
             * We haven't yet created a project that corresponds to
             * this business unit; create it now.
             */
            run('shares project ' + params.unit);
            set('mountpoint', '/export/' + params.unit);
            run('commit');
            run('shares select ' + params.unit);
        }

        run('filesystem ' + params.name);
        run('commit');
        return ('Created new share "' + params.name + '"');
    }
};
```

Workflow Execution Auditing and Reporting

Workflows may emit audit records by calling the `audit()` function. The `audit` function's only argument is a string that is to be placed into the audit log.

Using the `audit()` function shows the actual user who executed the workflow only if `setid` is set to `false`. However, if a workflow is owned by `root` and `setid` is set to `true`, audit logs will show `root` as the user, even if the workflow was run by another user.

To determine the user that is executing the workflow regardless of what `setid` is set to, use the `whoami()` function.

Example 9-7   Workflow Testing `whoami` Function

```
   var workflow = {
        name: "Test whoami",
        description: "Print current username",
        execute: function () {
                return ("Hello " + whoami());
        }
};
```

For complicated workflows that may require some time to execute, it can be useful to provide clear progress to the user executing the workflow. To allow the execution of a workflow to be reported in this way, the `execute` member should return an array of *steps*. Each array element must contain the following members:

Table 9-6   Required Members for Execution Reporting

Required Member	Type	Description
`step`	String	String that denotes the name of the execution step
`execute`	Function	Function that executes the step of the workflow

As with the `execute` function on the workflow as a whole, the `execute` member of each step takes as its argument an object that contains the parameters to the workflow.

Example 9-8   Workflow Execution Reporting

As an example, the following is a workflow that creates a new project, share, and audit record over three steps:

```
var steps = [ {
    step: 'Checking for associated project',
    execute: function (params) {
        try {
            run('shares select ' + params.unit);
        } catch (err) {
            if (err.code != EAKSH_ENTITY_BADSELECT)
                throw (err);

            /*
             * We haven't yet created a project that corresponds to
             * this business unit; create it now.
```

```
                 */
                run('shares project ' + params.unit);
                set('mountpoint', '/export/' + params.unit);
                run('commit');
                run('shares select ' + params.unit);
            }
        }
    }, {
        step: 'Creating share',
        execute: function (params) {
            run('filesystem ' + params.name);
            run('commit');
        }
    }, {
        step: 'Creating audit record',
        execute: function (params) {
            audit('created "' + params.name + '" in "' + params.unit);
        }
    } ];

    var workflow = {
        name: 'Create share',
        description: 'Creates a new share in a business unit',
        parameters: {
            name: {
                label: 'Name of new share',
                type: 'String'
            },
            unit: {
                label: 'Business unit',
                type: 'ChooseOne',
                options: [ 'development', 'finance', 'qa', 'sales' ],
                optionlabels: [ 'Development', 'Finance',
                    'Quality Assurance', 'Sales/Administrative' ],
            }
        },
        validate: function (params) {
            try {
                run('shares select ' + params.unit);
                run('select ' + params.name);
            } catch (err) {
                if (err.code == EAKSH_ENTITY_BADSELECT)
                    return;
            }

            return ({ name: 'share already exists' });
        },
        execute: function (params) { return (steps); }
    };
```

Using the `mail` function, workflows can deliver certain outputs of the workflow via email. The `mail` function must contain the following arguments: an object with `to` and `subject`, and a `messageBody` string.

Example 9-9 Workflow Execution with a Mailer

```
    var workflow = {
       name: 'email controller state',
       description: 'email controller state',
       execute: function () {

           // verify state of the controller
```

```
                var faulted = run('maintenance hardware "chassis-000" get faulted');

                var messageBody = faulted;

                emailAddress = 'first.last@xyz.com';
                subjectLine = 'Controller State';
                mail({To: emailAddress, Subject: subjectLine}, messageBody);

        }
};
```

Understanding Workflow Versioning

There are two aspects of versioning with respect to workflows: the first is the expression of the version of the Oracle ZFS Storage Appliance software that the workflow depends on, and the second is the expression of the version of the workflow itself. Versioning is expressed through two optional members to the workflow:

Table 9-7 Optional Members for Versioning

Optional Member	Type	Description
`required`	String	The minimum version of the appliance software required to run this workflow, including the minimum year, month, day, build and branch
`version`	String	Version of this workflow, in dotted decimal (major.minor.micro) form

Appliance Versioning - To express a minimally required version of the Oracle ZFS Storage Appliance software, add the optional `required` field to your workflow. The appliance is versioned in terms of the year, month and day on which the software was built, followed by a build number and then a branch number, expressed as `year.month.day.build-branch`. For example `2018.04.10,12-0` would be the twelfth build of the software originally build on April 10th, 2018. To get the version of the current appliance kit software, run the `configuration version get version` CLI command, or look at the **Version** field in the **System** screen in the BUI. Here's an example of using the `required` field:

Example 9-10 Using the Workflow `required` Field

Here's an example of using the `required` field:

```
var workflow = {
name: 'Configure FC',
description: 'Configures fibre channel target groups',
        required: '2018.12.25,1-0',
        ...
```

If a workflow requires a version of software that is newer than the version loaded on the appliance, the attempt to upload the workflow will fail with a message explaining the mismatch.

Workflow Versioning - In addition to specifying the required version of the appliance software, workflows themselves may be versioned with the `version` field. This string denotes the major, minor and micro numbers of the workflow version, and allows multiple versions of the same workflow to exist on the machine. When uploading a workflow, any *compatible*, *older* versions of the same workflow are deleted. A workflow is deemed to be *compatible* if it has the same major number, and a workflow is considered to be *older* if it has a lower version number.

Therefore, uploading a workflow with a version of `2.1` will remove the same workflow with version `2.0` (or version `2.0.1`) but not `1.2` or `0.1`.

Using Workflows for Alert Actions

A workflow can be executed as an alert action. A workflow that is executed as an alert action is the selected **Workflow** for the **Execute workflow Alert action** (BUI), or is the value of the `workflow` property of the `execute_workflow` handler (CLI), as described in Adding Alert Actions.

To enable a workflow to be executed as an alert action, the `workflow` global variable must specify an `alert: true` member.

By default, workflows execute as the user that is executing the workflow. However, a workflow that is executed as an alert action executes by default as the user that created the workflow. To be able to be executed as an alert action, the `workflow` global variable of the workflow must specify a `setid: true` member.

Alert actions have a single object parameter that has the following members.

Table 9-8 Required Members for Alert Execution Context

Required Member	Type	Description
`class`	String	The class of the alert
`code`	String	The code of the alert
`items`	Object	An object describing the alert
`timestamp`	Date	Time of alert

The `items` member of the parameters object has the following members:

Table 9-9 Required Members for the `items` Member

Required Member	Type	Description
`url`	String	The URL of the web page describing the alert
`action`	String	The action that should be taken by the user in response to the alert
`impact`	String	The impact of the event that precipitated the alert
`description`	String	A human-readable string describing the alert
`severity`	String	The severity of the event that precipitated the alert

Example 9-11 Workflow Auditing Failure to Reboot

Workflows that execute as alert actions can use the `audit` function to generate audit log entries. For example, debugging information should be written to the audit log.

See the multiple uses of the `audit` function in the following example. This workflow:

- Is executed in response to alert `params.uuid`, specifying this workflow to execute.
- Reboots the system only if both controllers are in the clustered state.
- Audits any failure to reboot.

```
var workflow = {
       name: 'Failover',
       description: 'Fail the node over to its clustered peer',
       alert: true,
       setid: true,
       execute: function (params) {
               var uuid = params.uuid;
               var clustered = 'AKCS_CLUSTERED';

               audit('attempting failover in response to alert ' + uuid);

               try {
                       run('configuration cluster');
               } catch (err) {
                       audit('could not get clustered state; aborting');
                       return;
               }

               if ((state = get('state')) != clustered) {
                       audit('state is ' + state + '; aborting');
                       return;
               }

               if ((state = get('peer_state')) != clustered) {
                       audit('peer state is ' + state + '; aborting');
                       return;
               }

               run('cd /');
               run('confirm maintenance system reboot');
       }
};
```

Creating and Posting Custom Alerts from Within a Workflow

An alert can be posted from within a workflow in response to an event that is defined in the workflow. The alert action can be created in the workflow, or by using the BUI or CLI. The alert must be posted from within the workflow.

Custom alerts can be used to help enforce administrative policy or compliance. Custom alerts can also help diagnose problems with a workflow.

The following table describes the authorizations that a user must have to create custom alert actions and post custom alerts.

Table 9-10 Authorizations Required to Use Custom Alerts

Task	BUI Authorization	CLI Authorization	Description
Create a custom alert action	Scope: Alerts Authorization: configure	Scope: `alert` Authorization: `allow_configure`	Required to create a custom alert action by using the BUI, the CLI, or the `createalert` function in a workflow
Post a custom alert	Scope: Alerts Authorization: post	Scope: `alert` Authorization: `allow_post`	Required to use the `postalert` function in a workflow
Execute a workflow	Scope: Workflow Authorization: read	Scope: `workflow` Authorization: `allow_read`	Required to execute a workflow

For instructions for granting authorizations to users, see Configuring Users.

Creating a Custom Alert

Use one of the following methods to create a custom alert:

- Specify the **Custom** or `custom` event category when creating an alert action, as described in the following documentation:
  - Adding an Alert Action for an Event Defined in a Workflow - BUI, CLI
  - Custom Alerts in *Oracle ZFS Storage Appliance RESTful API Guide, Release OS8.8.x*
- Use the `createalert` function in a script or workflow. To use `createalert` in a script, see Using the Custom Alert Functions. The remainder of this section describes how to use `createalert` in a workflow.

> **Note:**
>
> Each time `createalert` is called, a new alert is created with a different UUID. Rather than create more copies of the same alert action each time a workflow that uses `createalert` is executed, you might want to use the BUI or CLI to create a custom alert action, and then pass the UUID of that alert action to the `postalert` function, as shown in example "Posting a Custom Alert by Using an Existing Alert Action UUID" in Posting a Custom Alert.

The `createalert` function takes the following parameters, and returns the UUID of the custom alert action that is created.

Table 9-11 Parameters of the `createalert` Function

Parameter	Type	Description
`actions`	Object	Required. A list of handlers (alert actions) along with any arguments. See the "CLI Action Type" column of table "Alert Action Types" in Alert Action Types. **Note:** If `execute_workflow` is specified as an action (`handler`), the executed workflow cannot post an alert.
`severity`	String	Optional. The severity of the event that precipitated the alert. Valid values are: `Critical`, `Minor`, or `Major`.
`description`	String	Required. A description of the event that precipitated the alert.
`response`	String	Optional. A description of the actions that will be performed by the system to mitigate the effects of this event.
`impact`	String	Optional. A description of the effect that this event has on the appliance.
`recommended_action`	String	Optional. A description of the actions that the administrator should take to mitigate the effects of this event.

Example 9-12 Creating a Custom Alert from Within a Workflow

This example shows a workflow that creates a custom alert with UUID `custom_alert_uuid`.

```
var workflow = {
    name:        'createalert',
    description: 'Create a Custom Alert',
    version:     '1.0',
    origin:      'Oracle',
    alert:       false,
    setid:       true,
    execute:     function () {
                     var actions = [{
                         handler: 'email',
                         args: {
                             address: 'admin@example.com',
                             subject: 'Custom Alert Response'
                         }
                     }];
                     var createparams = {
                         description: 'createalert from within a workflow'
                     };
                     var custom_alert_uuid = createalert(actions, createparams);
                 }
};
```

Creating a custom alert by using a workflow causes the alert to appear in the alert actions list in the BUI and CLI.

Creating a custom alert by using a workflow creates an audit log entry with the following summary:

```
Workflow name_of_workflow: created custom alert value_of_custom_alert_uuid
```

Creating a custom alert by using the BUI or CLI or by running a script with the `script` command creates an audit log entry with the following summary:

```
Created custom alert value_of_custom_alert_uuid
```

Posting a Custom Alert

Use the `postalert` function to post a custom alert from within a workflow in response to an event that occurs in that workflow.

While the `createalert` function can be called from within a script by using the `script` command, the `postalert` function can only be called from within a workflow definition.

The `postalert` function takes the same parameters that the `createalert` function takes except for the first parameter: `createalert` takes a list of handlers or alert actions, and `postalert` takes the UUID of the alert to be posted. For descriptions of the other parameters, see table "Parameters of the `createalert` Function" in Creating a Custom Alert.

Values of the parameters that are optional for `createalert` (`severity`, `response`, `impact`, and `action`) are not optional for `postalert`. For the `postalert` call, these values are determined according to the following rules:

- If values for optional parameters are provided to `createalert`, but not to the `postalert` call for that UUID, then the `postalert` call inherits those parameter values from the corresponding `createalert` call.

- If values for optional parameters are not provided to createalert, but are provided to the postalert call for that UUID, then the postalert call uses those values specified in the postalert call.
- If values for optional parameters are provided to both createalert and the corresponding postalert, then each call uses the parameter values specified in that call.
- If values for optional parameters are not provided to neither createalert nor the corresponding postalert, then an error message tells the user to provide values to the postalert call.

The postalert function returns the UUID of the custom alert that is posted.

Example 9-13 Creating and Posting a Custom Alert from Within a Workflow

This example posts the alert that was created in the previous example. The postalert call is in response to an event that occurred in the workflow. In this example, the code that defines the event that occurs prior to the postalert call is omitted.

```
var workflow = {
    name:        'createalert and postalert',
    description: 'Create and Post a Custom Alert',
    version:     '1.0',
    origin:      'Oracle',
    alert:       false,
    setid:       true,
    execute:     function () {
                     var actions = [{
                         'handler': 'resume_dataset',
                         'args': {
                             'dataset': 'dataset_to_resume'
                         }
                     }];
                     var createparams = {
                         description: 'createalert and postalert from within a workflow'
                     };
                     var postparams = {
                         severity: 'Minor',
                         description: 'postalert from within a workflow',
                         response: 'The alert action resumes dataset dataset_to_resume',
                         impact: 'What happened to the appliance',
                         recommended_action: 'What the administrator should do'
                     };
                     var custom_alert_uuid = createalert(actions, createparams);
                     var posted_alert_uuid = postalert(custom_alert_uuid, postparams);
                 }
};
```

Executing a workflow that calls the postalert function creates an alert log entry with the following summary:

```
Custom: name_of_workflow
```

Example 9-14 Posting a Custom Alert by Using an Existing Alert Action UUID

Rather than proliferate copies of the same alert action, you might want to use the UUID of an existing custom alert as the first argument of the postalert function.

1. Use the BUI, CLI, a script, a workflow, or RESTful API to create a custom alert action.
2. Use the BUI, CLI, or audit log to retrieve the UUID of the custom alert action.
3. Use the retrieved UUID as the first argument to the postalert function.

```
var workflow = {
    name:        'postalert',
    description: 'Post a Custom Alert using existing uuid',
    version:     '1.0',
    origin:      'Oracle',
    alert:       false,
    setid:       true,
    execute:     function () {
                     var postparams = {
                         severity: 'Minor',
                         description: 'postalert from within a workflow',
                         response: 'What the system will do',
                         impact: 'What happened to the appliance',
                         recommended_action: 'What the administrator should do'
                     };
                     var posted_alert_uuid = postalert('uuid_of_existing_custom_alert',
postparams);
                 }
};
```

Using Scheduled Workflows

A workflow can be started via a timer event by defining a schedule for the workflow. To create and maintain schedules for a workflow, the property `scheduled` must be added to the Object workflow and must be set to `true`. See the `scheduled` property in example "Using Workflow Properties" in Coding Workflow Schedules.

Schedules can be created in the following ways:

- Using the CLI interface as described in Using a Scheduled Workflow
- Using an array type property named `schedules` in the Object workflow as described in Coding Workflow Schedules

Schedules can be changed via the CLI interface by an administrator with appropriate authorizations for the workflow.

Using a Scheduled Workflow

Once a workflow that contains the `scheduled: true` property is loaded onto an Oracle ZFS Storage Appliance system, the schedule can be defined or modified via the CLI interface. See the `scheduled` property in the example for "Using Workflow Properties" in Coding Workflow Schedules.

Each schedule entry consists of the following properties:

Table 9-12 Workflow Schedule Properties

Property	Type	Description
`NAME`	String	Name of the schedule. This value is system generated.
`frequency`	String	Values can be `minute`, `halfhour`, `hour`, `day`, `week`, `month`.
`day`	String	Values can be `Monday`, `Tuesday`, `Wednesday`, `Thursday`, `Friday`, `Saturday`, or `Sunday`. This value can be set only when `frequency` is set to `week` or `month`.
`hour`	String	Values can be `00, 01,…, 23`. This value can be set only when `frequency` is set to `day`, `week`, or `month`.

Table 9-12 (Cont.) Workflow Schedule Properties

Property	Type	Description
`minute`	String	`00, 01,…, 59`. This value can be set only when `hour` is set.

Example 9-15 Scheduled Workflow in the CLI

This example shows how to add two execution times to a workflow named "My Scheduled Workflow". The first execution time is daily at 10:05 a.m. The second scheduled execution time for this workflow is each Monday at 1:15 p.m.

```
hostname:> maintenance workflows
hostname:maintenance workflows> select name="My Scheduled Workflow"
hostname:maintenance workflow-002> schedules
hostname:maintenance workflow-002 schedules> create
hostname:maintenance workflow-002 schedule (uncommitted)> set frequency=day
                    frequency = day (uncommitted)
hostname:maintenance workflow-002 schedule (uncommitted)> set hour=10
                         hour = 10 (uncommitted)
hostname:maintenance workflow-002 schedule (uncommitted)> set minute=05
                       minute = 05 (uncommitted)
hostname:maintenance workflow-002 schedule (uncommitted)> commit
hostname:maintenance workflow-002 schedules> list
NAME                 FREQUENCY            DAY                  HH:MM
schedule-001         day                  -                    10:05
hostname:maintenance workflow-002 schedules> create
hostname:maintenance workflow-002 schedule (uncommitted)> set frequency=week
                    frequency = week (uncommitted)
hostname:maintenance workflow-002 schedule (uncommitted)> set day=Monday
                           day = Monday (uncommitted)
hostname:maintenance workflow-002 schedule (uncommitted)> set hour=13
                          hour = 13 (uncommitted)
hostname:maintenance workflow-002 schedule (uncommitted)> set minute=15
                       minute = 15 (uncommitted)
hostname:maintenance workflow-002 schedule (uncommitted)> commit
hostname:maintenance workflow-002 schedules> list
NAME                 FREQUENCY            DAY                  HH:MM
schedule-001         day                  -                    10:05
schedule-002         week                 Monday               13:15
hostname:maintenance workflow-002 schedules>
```

Coding Workflow Schedules

As an alternative to defining schedules by using the CLI interface as described in Using a Scheduled Workflow, schedules can be defined in the workflow code as a property in the Object workflow. The properties used in Object workflows are different from the properties used in CLI schedule creation.

The following properties are used in Object workflows:

Table 9-13 Workflow Schedule Properties

Property	Type	Description
`offset`	Number	Determines the starting point in the defined period

Table 9-13 (Cont.) Workflow Schedule Properties

Property	Type	Description
`period`	Number	Defines the frequency of the schedule
`unit`	String	Specifies if either `seconds` or `month` are used as the unit in the offset and period definition

The schedules that are coded into the workflow can be changed by an administrator with appropriate authorizations to the workflow as shown in Using a Scheduled Workflow.

Example 9-16 Using Workflow Properties

The following code example illustrates the use of the properties. Note that inline arithmetic helps to make the offset and period declarations more readable.

```
// Example of using Schedule definitions within a workflow
var MyTextObject = {
    MyVersion:  '1.0',
    MyName:     'Example 9',
    MyDescription:  'Example of use of Timer',
    Origin:     'Oracle'
};

var MySchedules = [
    // half hr interval
    { offset: 0, period: 1800, units: "seconds" },
    // offset 2 days, 4hr, 30min , week interval
    {offset: 2*24*60*60+4*60*60+30*60, period: 604800,units: "seconds" }
];

var workflow = {
    name:       MyTextObject.MyName,
    description:    MyTextObject.MyDescription,
    version:    MyTextObject.MyVersion,
    alert:      false,
    setid:      true,
    schedules:  MySchedules,
    scheduled:  true,
    origin:     MyTextObject.Origin,
    execute:    function () {
                audit('workflow started for timer');
               }
}
```

The property units in the Object `MySchedules` specifies the type of units used for the properties `offset` and `period`. They can be set to either `seconds` or `month`. The property `period` specifies the frequency of the event, and the `offset` specifies the units within the period. In the previous example, the period in the second schedule is set for a week, starting at the second day, at 4:30 a.m. Multiple schedules can be defined in the property schedules.

The Object `MySchedules` in the example uses the following three properties:

- `offset` - This is the starting offset from January 1, 1970 for the schedule. The offset is given in the units defined by the property `units`.
- `period` - This is the period between recurrences of the schedule, which is also given in the units defined by the property `units`.
- `units` - This can be defined in `seconds` or `months`.

The starting point for weekly schedules is Thursday. This is because the epoch is defined as starting on 1 Jan 1970, which was a Thursday.

Example 9-17   Workflow Schedule Shown in the CLI

In the previous example, the period in the second schedule uses a starting offset of 2 days + 4 hours + 30 minutes. This results in the starting date being January 3, 1970 at 4:30 a.m. The schedule recurs weekly indefinitely every Saturday at 4:30 a.m. Below, you can see the display of the schedule in the CLI.

```
hostname:> maintenance workflows
hostname:maintenance workflows> list
WORKFLOW     NAME                           OWNER SETID ORIGIN        VERSION
workflow-018 Check metaslab_unload_delay    root  false Oracle ZFSSA  1.0
workflow-019 Check metaslab_unload_timeout  root  false Oracle ZFSSA  1.0
workflow-020 Example 9                      root  true  <local>       1.0
workflow-021 Set DNLC size and ncsize       root  false Data and App  2.0
workflow-022 Stop Existing SSH Sessions     root  false <local>       0.2
hostname: maintenance workflows> select name="Example 9"
hostname: maintenance workflow-020> schedules
hostname: maintenance workflow-020 schedules> ls
Schedules:
NAME                 FREQUENCY            DAY                  HH:MM
schedule-000         halfhour             -                    --:00
schedule-001         week                 Saturday             04:30
```

Creating a Worksheet Based on a Specified Drive Type

Here is an example workflow that creates a worksheet based on a specified drive type:

Example 9-18   Workflow Device Type Selection

```
var steps = [ {
    step: 'Checking for existing worksheet',
    execute: function (params) {
        /*
         * In this step, we're going to see if the worksheet that
         * we're going to create already exists.  If the worksheet
         * already exists, we blow it away if the user has indicated
         * that they desire this behavior.  Note that we store our
         * derived worksheet name with the parameters, even though
         * it is not a parameter per se; this is explicitly allowed,
         * and it allows us to build state in one step that is
         * processed in another without requiring additional global
         * variables.
         */
        params.worksheet = 'Drilling down on ' + params.type + ' disks';

        try {
            run('analytics worksheets select name="' +
                params.worksheet + '"');

            if (params.overwrite) {
                run('confirm destroy');
                return;
            }

            throw ('Worksheet called "' + params.worksheet +
                '" already exists!');
        } catch (err) {
            if (err.code != EAKSH_ENTITY_BADSELECT)
```

```
                       throw (err);
               }
           }
        }, {
           step: 'Finding disks of specified type',
           execute: function (params) {
               /*
                * In this step, we will iterate over all chassis, and for
                * each chassis iterates over all disks in the chassis,
                * looking for disks that match the specified type.
                */
               var chassis, name, disks;
               var i, j;

               run('cd /');
               run('maintenance hardware');

               chassis = list();
               params.disks = [];

               for (i = 0; i < chassis.length; i++) {
                   run('select ' + chassis[i]);

                   name = get('name');
                   run('select disk');
                   disks = list();

                   for (j = 0; j < disks.length; j++) {
                       run('select ' + disks[j]);

                       if (get('use') == params.type) {
                           params.disks.push(name + '/' +
                               get('label'));
                       }

                       run('cd ..');
                   }

                   run('cd ../..');
               }

               if (params.disks.length === 0)
                   throw ('No ' + params.type + ' disks found');
               run('cd /');
           }
        }, {
           step: 'Creating worksheet',
           execute: function (params) {
               /*
                * In this step, we're ready to actually create the worksheet
                * itself:  we have the disks of the specified type and
                * we know that we can create the worksheet.  Note that we
                * create several datasets:  first, I/O bytes broken down
                * by disk, with each disk of the specified type highlighted
                * as a drilldown.  Then, we create a separate dataset for
                * each disk of the specified type.  Finally, note that we
                * aren't saving the datasets -- we'll let the user do that
                * from the created worksheet if they so desire.  (It would
                * be straightforward to add a boolean parameter to this
                * workflow that allows that last behavior to be optionally
                * changed.)
                */
               var disks = [], i;

               run('analytics worksheets');
               run('create "' + params.worksheet + '"');
               run('select name="' + params.worksheet + '"');
               run('dataset');
               run('set name=io.bytes[disk]');

               for (i = 0; i < params.disks.length; i++)
                   disks.push('"' + params.disks[i] + '"');

               run('set drilldowns=' + disks.join(','));
               run('commit');

               for (i = 0; i < params.disks.length; i++) {
                   run('dataset');
                   run('set name="io.bytes[disk=' +
                       params.disks[i] + ']"');
                   run('commit');
               }
           }
       } ];

       var workflow = {
           name: 'Disk drilldown',
           description: 'Creates a worksheet that drills down on system, ' +
               'cache, or log devices',
           parameters: {
               type: {
                   label: 'Create a new worksheet drilling down on',
                   type: 'ChooseOne',
                   options: [ 'cache', 'log', 'system' ],
                   optionlabels: [ 'Cache', 'Log', 'System' ]
               },
               overwrite: {
                   label: 'Overwrite the worksheet if it exists',
                   type: 'Boolean'
               }
           },
           execute: function (params) { return (steps); }
       };
```

Uploading and Executing Workflows Using the BUI

You can upload a workflow to Oracle ZFS Storage Appliance by clicking on the plus icon ⊕ . You can execute a workflow by clicking on the row that specifies the workflow, or by hovering over the workflow row and clicking the execute icon ➲ . You can also edit or delete a workflow by hovering over the workflow row, and clicking the appropriate icon.

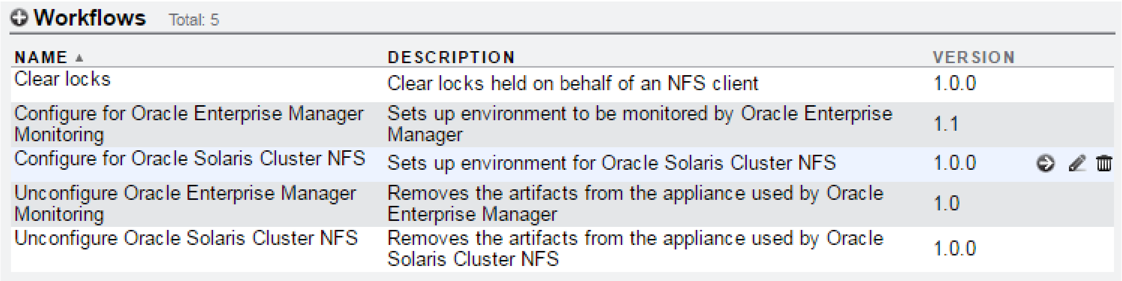

NAME ▲	DESCRIPTION	VERSION
Clear locks	Clear locks held on behalf of an NFS client	1.0.0
Configure for Oracle Enterprise Manager Monitoring	Sets up environment to be monitored by Oracle Enterprise Manager	1.1
Configure for Oracle Solaris Cluster NFS	Sets up environment for Oracle Solaris Cluster NFS	1.0.0
Unconfigure Oracle Enterprise Manager Monitoring	Removes the artifacts from the appliance used by Oracle Enterprise Manager	1.0
Unconfigure Oracle Solaris Cluster NFS	Removes the artifacts from the appliance used by Oracle Solaris Cluster NFS	1.0.0

To view the list of expanded workflows, hold **Shift** and click the plus icon ⊕ . To hide the expanded list, hold **Shift** and click the plus icon ⊕ again.

Downloading Workflows Using the CLI

Use the following procedure to download workflows using the CLI.

1. Workflows are downloaded to Oracle ZFS Storage Appliance via the `download` command, which is similar to the mechanism used for software updates:

   ```
   hostname:maintenance workflows> download
   hostname:maintenance workflows download (uncommitted)> get
                             url = (unset)
                            user = (unset)
                        password = (unset)
   ```

2. You must set the `url` property to be a valid URL for the workflow. This may be either local to your network or over the internet. The URL can be either HTTP (beginning with "http://") or FTP (beginning with "ftp://"). If user authentication is required, it may be a part of the URL (for example, "ftp://myusername:mypasswd@myserver/export/foo"), or you may leave the username and password out of the URL and instead set the `user` and `password` properties.

   ```
   hostname:maintenance workflows download (uncommitted)> set url=ftp://foo/
   example1.akwf
                             url = ftp://foo/example1.akwf
   hostname:maintenance workflows download (uncommitted)> set user=bmc
                            user = bmc
   hostname:maintenance workflows download (uncommitted)> set password
   Enter password:
                        password = (set)
   hostname:maintenance workflows download (uncommitted)> commit
   Transferred 138 of 138 (100%) ... done
   ```

Listing Workflows Using the CLI

Use the following procedure to list workflows using the CLI.

1. To list workflows, use the `list` command from the `maintenance workflows` context:

   ```
   hostname:maintenance workflows> list
   WORKFLOW     NAME                           OWNER SETID ORIGIN        VERSION
   workflow-018 Check metaslab_unload_delay    root  false Oracle ZFSSA  1.0
   workflow-019 Check metaslab_unload_timeout  root  false Oracle ZFSSA  1.0
   workflow-020 Example 9                      root  true  <local>       1.0
   workflow-021 Set DNLC size and ncsize       root  false Data and App  2.0
   workflow-022 Stop Existing SSH Sessions     root  false <local>       0.2
   ```

2. To select a workflow, use the `select` command:

   ```
   hostname:maintenance workflows> select workflow-018
   hostname:maintenance workflow-018>
   ```

3. To get a workflow's properties, use the `get` command from within the context of the selected workflow:

   ```
   hostname:maintenance workflow-018> get
                             name = Check metaslab_unload_delay
                      description = Check the value of the metaslab_unload_delay
   ```

```
property
                          uuid = uuid
                      checksum = checksum
                   installdate = 2019-5-8 20:54:33
                         owner = root
                        origin = Oracle ZFSSA
                         setid = false
                         alert = false
                       version = 1.0
                     scheduled = true
```

Executing Workflows Using the CLI

Use the following procedure to execute workflows using the CLI.

1. To execute a workflow, use the `run` command from within the context of the selected workflow.

   ```
   hostname:maintenance workflow-000> run
   ```

2. The context will become a captive context in which parameters must be specified for a workflow that has parameters:

   ```
   hostname:maintenance workflow-000> run
   hostname:maintenance workflow-000 run(uncommitted)> get
                            type = (unset)
                       overwrite = (unset)
   ```

   For a workflow with no parameters, you can `commit` directly after entering the captive context:

   ```
   hostname:maintenance workflow-000> run
   hostname:maintenance workflow-000 run(uncommitted)> commit
   hello world!
   ```

3. Any attempt to `commit` the execution of the workflow without first setting the requisite parameters will result in an explicit failure:

   ```
   hostname:maintenance workflow-000 run(uncommitted)> commit
   error: cannot execute workflow without setting property "type"
   ```

4. To execute the workflow, set the specified parameters, and then use the `commit` command:

   ```
   hostname:maintenance workflow-000 run (uncommitted)> set type=system
                            type = system
   hostname:maintenance workflow-000 run (uncommitted)> set overwrite=true
                       overwrite = true
   hostname:maintenance workflow-000 run (uncommitted)> commit
   ```

5. If the workflow has specified steps, those steps will be displayed via the CLI, for example:

   ```
   hostname:maintenance workflow-000 run (uncommitted)> commit
   Checking for existing worksheet ... done
   Finding disks of specified type ... done
   Creating worksheet ... done
   ```

Auditing Workflows Using the CLI

All workflows have a checksum property computed by the system. This checksum is the SHA-256 digest of the workflow content. To determine if a workflow has changed, compare the workflow checksum against its previous checksum.

To obtain the checksum property of a workflow, use the `get checksum` command from the `maintenance workflows` context:

```
hostname:maintenance workflows> select workflow-001
hostname:maintenance workflow-001> get checksum
checksum = 15f4188643d7add37b5ad8bda6d9b4e7210f1cd890a73df176382e800aec
```

10 Integration

Oracle ZFS Storage Appliance is engineered to seamlessly integrate with other Oracle products as well as third-party and virtualized environments. The appliance provides a full suite of data protocols to communicate with a wide variety of application hosts. To improve application performance, provide effective backup support, or more tightly integrate with your application environment and manage and monitor storage appliances, see these resources for a complete list of available downloadable plug-ins, and for technical briefs and documentation detailing configuration and deployment best practices and recommendations for maximizing performance:

- Oracle ZFS Storage Appliance Plug-in Downloads
- NAS Storage Documentation

Configuring Oracle ZFS Storage Appliance for Oracle Database Clients

Oracle ZFS Storage Appliance offers a number of unique features designed to integrate with Oracle Database clients, including Hybrid Columnar Compression (HCC) and Oracle Intelligent Storage Protocol (OISP).

To enable these features, the SNMP service on Oracle ZFS Storage Appliance must be configured to allow SNMP queries by database clients. The client uses this mechanism to identify a storage device as an Oracle ZFS Storage Appliance. For specifying the database client hostname or IP address as the trap destination, see Configuring SNMP to Serve Appliance Status (BUI).

To verify the appliance SNMP service is configured properly, run the `snmpget(1)` command from the client system, replacing <*host*> with the name or IP address of the appliance.

```
-bash-4.1$ snmpget -v1 -c public <host> .1.3.6.1.4.1.42.2.225.1.4.2.0
SNMPv2-SMI::enterprises.42.2.225.1.4.2.0 = STRING: "Oracle ZFS Storage Appliance"
```

Oracle Intelligent Storage Protocol

Oracle Intelligent Storage Protocol (OISP) enables the Oracle direct NFS (dNFS) client to encode and pass attributes associated with I/O requests to Oracle ZFS Storage Appliance. These attributes contain such information as the type of database file that the I/O request is targeting, the record size of the file, whether to cache the I/O data, and the identity of the database issuing the I/O request.

The appliance decodes these attributes, using them to simplify database configuration, increase database performance, and provide observability into the source of I/O workloads being generated by database clients.

Database Record Size

The Oracle dNFS client can pass the optimal record size based on the type of file for each I/O request. If a record size is passed, it overrides the `Database record size` property setting on the share or project. The record size can only be set for newly created files. If a file already exists, the record size is not changed.

Synchronous Write Bias Hint

The Oracle dNFS client can pass a write bias "hint" associated with write I/O requests that prompts the appliance to treat I/O requests as either latency sensitive or throughput oriented. If the hint is passed, it overrides the `Synchronous write bias` property setting on the share or project.

Analytics Breakdown by Database Name

The Oracle Database 12c or later dNFS client can pass the identification of the database (SID) or container database and pluggable database (SID:SID) responsible for issuing I/O requests. Oracle ZFS Storage Appliance analytics can display I/O statistics broken down by the SID name(s) of the database by selecting breakdown or drill by `Application ID`.

With OS8.7 and later firmware on Oracle ZFS Storage Appliance, additional OISP database analytics may be displayed. OISP operations by client, file name, database name, database file type, database function, share, project, size file offset, and latency are all available.

Caching Hints

The Oracle Database 12.2 or later dNFS client includes caching hints on I/O requests. Negative caching hints are included on I/O requests that do not expect to soon reference the data read or written again, such as datafile blocks read, and backup pieces written as part of Oracle Recovery Manager (Oracle RMAN) backup. This assists the appliance in making the best use of available memory in caching filesystem data. The main negatively cached operations are: Oracle RMAN reads and writes, Oracle Database datafile and redo log file creation, and Oracle Database Archiver reads and writes.

OISP-Capable Protocols and Clients

- Protocols: NFSv4.0 and NFSv4.1
- Clients: Oracle Database NFS (dNFS) client